Technische Universität Berlin
Fakultät für Elektrotechnik und Informatik
Lehrstuhl für Intelligente Netze
und Management Verteilter Systeme

arXiv:1511.01779v1 [cs.DC] 5 Nov 2015

On the Cost of Concurrency in Transactional Memory

vorgelegt von
Srivatsan Ravi (M.E.)

von der Fakultät IV – Elektrotechnik und Informatik
der Technischen Universität Berlin
zur Erlangung des akademischen Grades
Doktor der Ingenieurwissenschaften (Dr.-Ing.)
genehmigte Dissertation

Promotionsausschuss:

Vorsitzender:	Prof. Uwe Nestmann, Ph. D., TU Berlin
Gutachterin:	Prof. Anja Feldmann, Ph. D., TU Berlin
Gutachter:	Prof. Petr Kuznetsov, Ph. D., Télécom ParisTech
Gutachterin:	Prof. Hagit Attiya, Ph. D., The Technion
Gutachter:	Prof. Rachid Guerraoui, Ph. D., EPFL
Gutachter:	Prof. Michel Raynal, Ph. D., INRIA, Rennes

Tag der wissenschaftlichen Aussprache: 18 June, 2015

Berlin 2015
D 83

Eidesstattliche Erklärung

Ich versichere an Eides statt, dass ich diese Dissertation selbständig verfasst und nur die angegebenen Quellen und Hilfsmittel verwendet habe.

Datum Srivatsan Ravi (M.E.)

Abstract

Current general-purpose CPUs are *multicores*, offering multiple computing units within a single chip. The performance of programs on these architectures, however, does not necessarily increase proportionally with the number of cores. Designing concurrent programs to exploit these multicores emphasizes the need for achieving efficient *synchronization* among *threads* of computation. When there are several threads that *conflict* on the same data, the threads will need to coordinate their actions for ensuring correct program behaviour.

Traditional techniques for synchronization are based on *locking* that provides threads with exclusive access to shared data. *Coarse-grained* locking typically forces threads to access large amounts of data sequentially and, thus, does not fully exploit hardware concurrency. Program-specific *fine-grained* locking or *non-blocking* (*i.e.*, not using locks) synchronization, on the other hand, is a dark art to most programmers and trusted to the wisdom of a few computing experts. Thus, it is appealing to seek a middle ground between these two extremes: a synchronization mechanism that relieves the programmer of the overhead of reasoning about data conflicts that may arise from concurrent operations without severely limiting the program's performance. The *Transactional Memory (TM)* abstraction is proposed as such a mechanism: it intends to combine an easy-to-use programming interface with an efficient utilization of the concurrent-computing abilities provided by multicore architectures. TM allows the programmer to *speculatively* execute sequences of shared-memory operations as *atomic transactions* with *all-or-nothing* semantics: the transaction can either *commit*, in which case it appears as executed sequentially, or *abort*, in which case its update operations do not take effect. Thus, the programmer can design software having only sequential semantics in mind and let TM take care, at run-time, of resolving the conflicts in concurrent executions.

Intuitively, we want TMs to allow for as much *concurrency* as possible: in the absence of severe data conflicts, transactions should be able to progress in parallel. But what are the inherent costs associated with providing high degrees of concurrency in TMs? This is the central question of the thesis.

To address this question, we first focus on the *consistency* criteria that must be satisfied by a TM implementation. We precisely characterize what it means for a TM implementation to be *safe*, *i.e.*, to ensure that the view of *every* transaction could have been observed in some sequential execution. We then present several lower and upper bounds on the complexity of three classes of safe TMs: *blocking* TMs that allow transactions to delay or abort due to *overlapping* transactions, *non-blocking* TMs which adapt to *step contention* by ensuring that a transaction not encountering steps of overlapping transactions must commit, and *partially non-blocking* TMs that provide strong non-blocking guarantees (*wait-freedom*) to only a subset of transactions. We then propose a model for *hybrid* TM implementations that complement hardware transactions with software transactions. We prove that there is an inherent trade-off on the degree of concurrency allowed between hardware and software transactions and the costs introduced on the hardware. Finally, we show that *optimistic* synchronization techniques based on speculative executions are, in a precise sense, better equipped to exploit concurrency than inherently *pessimistic* techniques based on locking.

Zusammenfassung

Aktuelle Allzweck-CPUs haben *mehrere* Rechenkerne innerhalb eines einzelnen Chipsatzes. Allerdings erhöht sich die Leistung der Programme auf diesen Architekturen nicht notwendigerweise proportional in der Anzahl der Kerne. Das Entwerfen nebenläufiger Programme um diese Multicores zu nutzen, erfordert die Überwindung einiger nicht-trivialer Herausforderungen; die wichtigste ist, eine effiziente *Synchronisierung* der *Threads* der Berechnung herzustellen. Greifen mehrere Threads gleichzeitig auf dieselben Daten zu, müssen diese ihre Aktionen koordinieren, um ein korrektes Programmverhalten zu gewährleisten.

Die traditionelle Methode zur Synchronisierung ist "Locking", welches jeweils nur einem einzelnen Thread Zugriff auf gemeinsam genutzten Daten gewährt. Bei *grobkörnigem* "Lockingërfolgt der Zugang zu einer großen Menge von Daten meist seriell, sodass die Hardware-Parallelität nich in vollem Unfang ausgenutzt wird. Auf der anderen Seite stellt programmspezifisches *feinkörniges* Locking, oder auch nicht-blockierende (d.h. keine Locks benutzende) Synchronisierung, eine dunkle Kunst für die meisten Programmierer dar, welche auf die Weisheit weniger Computerexperten vertraut. So ist es angebracht, einen Mittelweg zwischen diesen beiden Extremen zu suchen: einen Synchronisierungsmechanismus, der den Programmierer bezüglich der *Datenkonflikte*, die aus gleichzeitigen Operationen entstehen, entlastet, ohne jedoch die Leistung des Programms zu stark zu beeinträchtigen. Die *Transactional Memory (TM)* Abstraktion wird als solcher Mechanismus vorgeschlagen: ihr Ziel ist es, eine einfach zu bedienende Programmierschnittstelle mit einer effizienten Nutzung der gleichzeitigen Computing-Fähigkeiten von Multicore-Architekturen zu kombinieren. TM erlaubt es dem Programmierer, Sequenzen von Operationen auf dem gemeinsamen Speicher als *atomare Transaktionen* mit *Alles-oder-Nichts* Semantik zu erklären: Die Transaktion wird entweder *übergeben*, und somit sequentiell ausgeführt, oder *abgebrochen*, sodass ihre Operationen nicht durchgeführt werden. Dies ermöglicht dem Programmierer, Software mit nur sequentieller Semantik zu konzipieren, und die aus gleichzeitger Ausführung entstehenden Konflikte TM zu überlassen.

Intuitiv sollen die TMs so viel Nebenläufigkeit wie möglich berücksichtigen: Falls keine Datenkonflikte vorhanden sind, sollen alle Transaktionen parallel ausgeführt werden. Gibt es in TMs Kosten, die durch diesen hohen Grad an *Nebenläufigkeit* entstehen? Das ist die zentrale Frage dieser Arbeit.

Um diese Frage zu beantworten, konzentrieren wir uns zunächst auf das Kriterium der *Konsistenz*, welche von der TM-Implementierung erfüllt werden muss. Wir charakterisieren auf präzise Art, was es für eine TM-Implementierung heißt, *sicher* zu sein, d.h. zu gewährleisten, dass die Sicht *einer jeden* Transaktion auch von einer sequentiellen Transaktion hätte beobachtet werden können. Danach präsentieren wir mehrere untere und obere Schranken für die Komplexität dreier Klassen von sicheren TMs: *blockierende* TMs, die Blockierungen oder Abbrüche der Transaktionen erlauben, sollten diese sich überlappen, *nicht-blockierende* TMs die einen schrittweisen Zugriffskonflikt berücksichtigen, d.h. Transaktionen, die keinen Zugriff überlappender anderer Transaktionen beobachten, müssen übergeben, und *partiell nicht-blockierende* TMs, die nur für eine Teilmenge von Transaktionen nicht-blockierend sind. Wir schlagen daraufhin ein Modell für *hybride* TM-Implementierungen vor, welches die Hardware Transaktionen mit Software Transaktionen ergänzt. Wir beweisen, dass es eine inherente Trade-Off zwischen Grad der erlaubten Nebenläufigkeit zwischen Hard- und Software Transaktionen und den Kosten der Hardware gibt. Schlussendlich beweisen wir, dass *optimistische*, auf *spekulativen* Ausführungen basierende, Synchronisierungstechniken, in einem präzisen Sinne, besser geeignet sind um Nebenläufigkeit auszunutzen als *pessimistische* Techniken, die auf "Locking"basieren.

Contents

1 Introduction **11**
- 1.1 Concurrency and synchronization . . . 11
  - 1.1.1 Concurrent computing overview . . . 11
  - 1.1.2 Synchronization using locks . . . 12
  - 1.1.3 Non-blocking synchronization . . . 13
- 1.2 Transactional Memory (TM) . . . 14
- 1.3 Summary of the thesis . . . 15
  - 1.3.1 Safety for transactional memory . . . 15
  - 1.3.2 Complexity of transactional memory . . . 16
  - 1.3.3 Hybrid transactional memory . . . 17
  - 1.3.4 Optimism for boosting concurrency . . . 18
- 1.4 Roadmap of the thesis . . . 19

2 Transactional Memory model **21**
- 2.1 TM interface and TM implementations . . . 21
- 2.2 TM-correctness . . . 24
- 2.3 TM-progress . . . 25
- 2.4 TM-liveness . . . 27
- 2.5 Invisible reads . . . 28
- 2.6 Disjoint-access parallelism (DAP) . . . 29
- 2.7 TM complexity metrics . . . 31

3 Safety for transactional memory **33**
- 3.1 Overview . . . 33
- 3.2 Safety properties . . . 34
- 3.3 Opacity and deferred-update(DU) semantics . . . 35
- 3.4 On the safety of du-opacity . . . 36
  - 3.4.1 Du-opacity is prefix-closed . . . 36
  - 3.4.2 The limit of du-opaque histories . . . 37
  - 3.4.3 Du-opacity is limit-closed for complete histories . . . 38
  - 3.4.4 Du-opacity vs. opacity . . . 41
- 3.5 Strict serializability with DU semantics . . . 43
- 3.6 Du-opacity vs. other deferred-update criteria . . . 43
  - 3.6.1 Virtual-world consistency . . . 44
  - 3.6.2 Transactional memory specification (TMS) . . . 45
- 3.7 Related work and Discussion . . . 48

4 Complexity bounds for blocking TMs **49**
- 4.1 Overview . . . 49
- 4.2 Sequential TMs . . . 50
  - 4.2.1 A quadratic lower bound on step complexity . . . 51
  - 4.2.2 Expensive synchronization in Transactional memory cannot be eliminated . . . 53
- 4.3 Progressive TMs . . . 53
  - 4.3.1 A linear lower bound on the amount of protected data . . . 54
  - 4.3.2 A constant stall and constant expensive synchronization strict DAP opaque TM . . . 55

4.4 Strongly progressive TMs . . . 61
4.4.1 A $\Omega(n \log n)$ lower bound on remote memory references . . . 61
4.4.2 A constant expensive synchronization opaque TM . . . 65
4.5 On the cost of permissive opaque TMs . . . 68
4.6 Related work and Discussion . . . 70

5 Complexity bounds for non-blocking TMs **73**
5.1 Overview . . . 73
5.2 Impossibility of weak DAP and invisible reads . . . 73
5.3 A linear lower bound on memory stall complexity . . . 76
5.4 A linear lower bound on expensive synchronization for RW DAP . . . 78
5.5 Algorithms for obstruction-free TMs . . . 81
5.5.1 An opaque RW DAP TM implementation . . . 82
5.5.2 An opaque weak DAP TM implementation . . . 86
5.6 Why Transactional memory should not be obstruction-free . . . 87
5.7 Related work and Discussion . . . 88

6 Lower bounds for partially non-blocking TMs **91**
6.1 Overview . . . 91
6.2 The space complexity of invisible reads . . . 92
6.3 Impossibility of strict DAP . . . 93
6.4 A linear lower bound on expensive synchronization for weak DAP . . . 95
6.5 Related work and Discussion . . . 98

7 Hybrid transactional memory (HyTM) **99**
7.1 Overview . . . 99
7.2 Modelling HyTM . . . 101
7.2.1 Direct and cached accesses . . . 101
7.2.2 Slow-path and fast-path transactions . . . 102
7.2.3 Instrumentation . . . 103
7.2.4 Impossibility of uninstrumented HyTMs . . . 104
7.3 A linear lower bound on instrumentation for progressive HyTMs . . . 106
7.4 Instrumentation-optimal progressive HyTM . . . 111
7.5 Providing partial concurrency at low cost . . . 116
7.6 Related work and Discussion . . . 117

8 Optimism for boosting concurrency **119**
8.1 Overview . . . 119
8.2 Concurrent implementations . . . 121
8.3 Locally serializable linearizability . . . 123
8.4 Pessimistic vs. optimistic synchronization . . . 125
8.4.1 Concurrency analysis . . . 126
8.4.2 Concurrency optimality . . . 127
8.5 Related work and Discussion . . . 132

9 Concluding remarks **133**

List of Figures **137**

List of Tables **139**

10 Bibliography **141**

Acknowledgements

In his wonderfully sarcastic critique of the scientific community in *His Master's Voice*, the great Polish writer Stanisław Lem refers to a *specialist* as a *barbarian whose ignorance is not well-rounded.* Writing a Ph. D. thesis is essentially an attempt at becoming a specialist on some topic; whether this thesis on Transactional Memory makes me one is a questionable claim, but I am culturedly not totally ignorant, I think. The thesis itself was a long time in the making and would not have been possible without the wonderful support and gratitudes I have received these past four years.

My advisors Anja Feldmann and Petr Kuznetsov guided me throughout my Ph. D. term.

A fair amount of whatever good I have learnt these past few years, both scientifically and meta-scientifically, I owe it to Petr. He taught me, by example, what it takes to achieve nontrivial scientific results. He spent several hours schooling me when I had misunderstood some topic and as such suffered the worst of my writing, especially in the first couple of years. He was hard on me when I did badly, but always happy for me when I did well. Apart from being a deep thinker and a brilliant researcher, his scientific integrity and mental discipline have indelibly made me a better human being and student of science. At a personal level, I wish to thank him and his family for undeserved kindness shown to me over the years.

I would like to thank Anja for the extraordinary amount of freedom she gave me to pursue my own research and the trust she placed in me, as she does in all her students.

I am especially grateful to Robbert Van Renesse and Bryan Ford, who gave me a taste for independent research and in many ways, helped shape the course of my graduate career.

I am of course extremely grateful to all my co-authors who allowed me to include content, written in conjunction with them, in the thesis. So special thanks to Dan Alistarh, Hagit Attiya, Vincent Gramoli, Sandeep Hans, Justin Kopinsky, Petr Kuznetsov and Nir Shavit.

The results in Chapter 3 were initiated during a memorable visit to the Technion, Haifa in the Spring of '12. I am thankful to have been hosted by Hagit Attiya and to have had the chance to work with her and Sandeep Hans. I am also very grateful to David Sainz for taking time off and introducing me to some beautiful parts of Israel.

Chapter 7 essentially stemmed from a visit to MIT in the summer of '13. I am very grateful to Dan Alistarh and Nir Shavit for hosting me. Special thanks to Justin Kopinsky, who helped keep our discussions alive during our lengthy dry spells when we were seemingly spending all our time thinking about the problem, but without producing any tangible results.

Chapter 8 represents the most excruciatingly painful part of the thesis purely in terms of the number of iterations the paper based on this chapter went through. Yet, it was a procedure from which I learnt a lot and I am very thankful to Vincent Gramoli, who initiated the topic during my visit to EPFL in Spring '11.

In general, I have benefitted immensely from just talking to researchers in distributed computing during conferences, workshops and research visits. These include Yehuda Afek, Panagiota Fatourou, Rachid Guerraoui, Maurice Herlihy, Victor Luchangco, Adam Morrison and Michel Raynal. Also, great thanks to the anonymous reviewers of my paper submissions whose critiques and comments helped improve the contents of the thesis.

Back here in Berlin, so many of my INET colleagues have shaped my thought processes and enriched my experience in grad school. Thanks to Stefan Schmid, whose ability to execute several tasks concurrently with minimal synchronization overhead, never ceases to amaze anyone in this group. I am also very grateful to Anja, Petr and Stefan for allowing me to be a Teaching Assistant in their respective courses. Apart from being great friends, Felix Poloczek and Matthias Rost have been wonderful office mates and indulged my random discussions about life and research. Arne Ludwig and Carlo Fürst have been great friends; Arne, thanks for all the football discussions and Carlo, for exposing me to some social life. Thanks to Dan Levin, who has been a great friend and always been there to motivate and give me a fillip whenever I needed it. Ingmar Poese has been a wonderful friend as well as a constant companion to the

movies. Franziska Lichtblau, Enric Pujol, Philipp Richter and others have been willing companions in ordering several late night dinners at the lab. Thomas Krenc and Philipp Schmidt were great co-TA's. Great thanks to all the other current and past members of INET: our system admin Rainer May, Marco Canini, Damien Foucard, Juhoon Kim, Gregor Schaffrath, Julius Schulz-Zander, Georgios Smaragdakis, Florian Streibelt, Lalith Suresh, Steve Uhlig and all the other members I have missed. Special thanks to our group secretaries, Birgit Hohmeier-Touré and Nadine Pissors, without whom there would be absolute chaos.

I would like to thank my flat mates of the last two years: Ingmar, Jennifer, Jose, Lily and Mathilda. Lastly, thanks to my family and friends outside of the academic sphere for tolerating me all these years.

1 Introduction

While the performance of programs would increase proportionally with the performance of a *singlecore* CPU, the performance of programs on *multicore* CPU architectures, however, does not necessarily increase proportionally with the number of cores. In order to exploit these multicores, the amount of concurrency provided by programs will need to increase as well. Designing concurrent programs that exploit the hardware concurrency provided by modern multicore CPU architectures requires achieving efficient *synchronization* among *threads* of computation. However, due to the *asynchrony* resulting from the CPU's context switching and scheduling policies, it is hard to specify reasonable bounds on relative thread speeds. This makes the design of efficient and correct concurrent programs a difficult task. The *Transactional Memory (TM)* abstraction [80, 117] is a synchronization mechanism proposed as a solution to this problem: it combines an easy-to-use programming interface with an efficient utilization of the concurrent-computing abilities provided by multicore architectures. This chapter introduces the TM abstraction and presents an overview of the thesis.

1.1 Concurrency and synchronization

In this section, we introduce the challenges of concurrent computing and overview the drawbacks associated with traditional synchronization techniques.

1.1.1 Concurrent computing overview

Shared memory model. A *process* represents a thread of computation that is provided with its own private memory which cannot be accessed by other processes. However, these independent processes will have to synchronize their actions in an asynchronous environment in order to *implement* a user application, which they do by communicating via the CPU's *shared memory*.

In the shared memory model of computation, processes communicate by *reading* and *writing* to a fragment of the shared memory, referred to as a *base object*, in a single *atomic* (*i.e.*, indivisible) instruction. Modern CPU architectures additionally allow processes to invoke certain powerful atomic *read-modify-write (rmw)* instructions [75], which allow processes to write to a base object subject to the check of an invariant. For example, the *compare-and-swap* instruction is a rmw instruction that is supported by most modern architectures: it takes as input $\langle old, new \rangle$ and atomically updates the value of a base object to *new* and returns *true iff* its value prior to applying the instruction is equal to *old*; otherwise it returns *false*.

Concurrent implementations. A *concurrent implementation* provides each process with an algorithm to apply CPU instructions on the shared base objects for the *operations* of the user application. For example, consider the problem of implementing a concurrent *list-based set* [81]. the *set* abstraction implemented as a *sorted linked list* supporting operations insert(v), remove(v) and contains(v); $v \in \mathbb{Z}$. The set abstraction stores a set of integer values, initially empty. The update operations, insert(v) and remove(v), return a boolean response, *true* if and only if v is absent (for insert(v)) or present (for remove(v)) in the list. After insert(v) is complete, v is present in the list, and after remove(v) is complete, v is absent in the list. The contains(v) returns a boolean, *true* if and only if v is present in the list. A concurrent implementation of the list-based set is simply an *emulation* of the set abstraction that is realized by processes applying the available CPU instructions on the underlying base objects.

Safety and liveness. What does it mean for a concurrent implementation to be correct? Firstly, the implementation must satisfy a *safety property*: there are no bad reachable states in any *execution* of the implementation. Intuitively, we characterize safety for a concurrent implementation of a data abstraction by verifying if the responses returned in the concurrent execution may have been observed in a *sequential execution* of the same. For example, the safety property for a concurrent list-based set implementation stipulates that the response of the set operations in a concurrent execution is consistent with some sequential execution of the list-based set. However, a concurrent set implementation that does not return any response trivially ensures safety; thus, the implementation must satisfy some *liveness property* specifying the conditions under which the operations must return a response. For example, one liveness property we may wish to impose on the concurrent list-based set is *wait-freedom*: every process completes the operations it invokes irrespective of the behaviour of other processes.

As another example, consider the *mutual exclusion* problem [41] which involves sharing some critical data resource among processes. The safety property for mutual exclusion stipulates that at most one process has access to the resource in any execution, in which case, we say that the process is inside the *critical section*. However, one may notice that an implementation which ensures that no process ever enters the critical section is trivially safe, but not very useful. Thus, the mutual exclusion implementation must satisfy some liveness property specifying the conditions under which the processes must eventually enter the critical section. For example, we expect that the implementation is *deadlock-free*: if every process is given the chance to execute its algorithm, some process will enter the critical section. In contrast to safety, a liveness property can be violated only in an infinite execution, *e.g.*, by no process ever entering the critical section.

In shared memory computing, we are concerned with deriving concurrent implementations with strong safety and liveness properties, thus emphasizing the need for efficient synchronization among processes.

1.1.2 Synchronization using locks

A *lock* is a concurrency abstraction that implements mutual exclusion and is the traditional solution for achieving synchronization among processes. Processes *acquire* a lock prior to executing code inside the critical section and *release* the lock afterwards, thereby allowing other processes to modify the data accessed by the code within the critical section. In essence, after acquiring the lock, the code within the critical section can be executed atomically. However, *lock-based* implementations suffer from some fundamental drawbacks.

Ease of designing lock-based programs. Ideally, to reduce the programmer's burden, we would like to take any sequential implementation and transform it to an efficient concurrent one with minimal effort. Consider a simple locking protocol that works for most applications: *coarse-grained* locking which typically serializes access to a large amount of data. Although trivial for the programmer to implement, it does not exploit hardware concurrency. In contrast, *fine-grained* locking may exploit concurrency better, but requires the programmer to have a good understanding of the data-flow relationships in the application and precisely specify which locks provide exclusive access to which data.

For example, consider the problem of implementing a concurrent list-based set. The *sequential* implementation of the list-based set uses a sorted linked list data structure in which each data item (except

the *tail* of the list) maintains a next field to provide a pointer to the successor. Every operation (insert, remove and contains) invoked with a parameter $v \in \mathbb{Z}$ traverses the list starting from the *head* up to the data item storing value $v' \geq v$. If $v' = v$, then contains returns *true*, remove(v) unlinks the corresponding element and returns true and insert(v) returns *false*. Otherwise, contains(v) and remove(v) return false, while insert(v) adds a new data item with value v to the list and returns *true*.

Given such a sequential implementation, we may derive a coarse-grained implementation of the list-based set by having processes acquire a lock on the head of the list, thus, forcing one operation to complete before the next starts. Alternatively, a fine-grained protocol may involve acquiring locks *hand-over-hand* [29]: a process holds the lock on at most two adjacent data items of the list. Yet, while such a protocol produces a correct set implementation [26], it is not a universal strategy that applies to other popular data abstractions like *queues* and *stacks*.

Composing lock-based programs. It is hard to compose smaller atomic operations based on locks to produce a larger atomic operation without affecting safety [81, 117]. Consider the *fifo queue* abstraction supporting the enqueue(v); $v \in \mathbb{Z}$ and dequeue operations. Suppose that we wish to solve the problem of atomically dequeuing from a queue Q_1 and enqueuing the item returned, say v, to a queue Q_2. While the individual actions of dequeuing from Q_1 and enqueuing v to Q_2 may be atomic, we wish to ensure that the combined action is atomic: no process must observe the absence of v or that it is present in both Q_1 and Q_2. A possible solution to this specific problem is to force a process attempting atomic modification of Q_1 and Q_2 to acquire a lock. Firstly, this requires prior knowledge of the identities of the two queue instances. Secondly, this solution does not exploit hardware concurrency since the lock itself becomes a concurrency bottleneck. Moreover, imagine that processes p_1 and p_2 need to acquire two locks L_1 and L_2 in order to atomically modify a set of queue instances. Without imposing a pre-agreed upon order on such lock acquisitions, there is the possibility of introducing *deadlocks* where processes wait infinitely long without completing their operations. For example, imagine the following concurrency scenario: process p_1 (and resp., p_2) holds the lock L_1 (and resp., L_2) and attempts to acquire the lock L_2 (and resp., L_1). Thus, process p_1 (and resp., p_2) waits infinitely long for p_2 (and resp., p_1) to complete its operation.

1.1.3 Non-blocking synchronization

It is impossible to derive lock-based implementations that provide *non-blocking* liveness guarantees, *i.e.*, some process completes its operations irrespective of the behaviour of other processes. In fact, even the weak non-blocking liveness property of *obstruction-freedom* [16] cannot be satisfied by lock-based implementations: a process must complete its operation if it eventually runs *solo* without interleaving events of other processes.

Concurrent implementations providing non-blocking liveness properties are appealing in practice since they overcome problems like deadlocks and *priority inversions* [30] inherent to lock-based implementations. Thus, non-blocking (without using locks) solutions using *conditional* rmw instructions like compare-and-swap have been proposed as an alternative to lock-based implementations. However, as with fine-grained locking, implementing correct non-blocking algorithms can be hard and requires hand-crafted problem-specific strategies. For example, the state-of-the-art list-based set implementation by Harris-Michael [70, 81, 105] is non-blocking: the insert and remove operations, as they traverse the list, *help* concurrent operations to physically remove data items (using compare-and-swap) that are logically deleted, *i.e.*, "marked for deletion". But one cannot employ an identical algorithmic technique for implementing a non-blocking queue [106], whose semantics is orthogonal to that of the set abstraction. Moreover, addressing the compositionality issue, as with lock-based solutions, requires ad-hoc strategies that are not easy to realize [81].

Algorithm 1.1 Sequential implementation of remove operation of list-based set

```
1:  remove(v):
2:    prev ← head
3:    curr ← read(prev.next)
4:    while (tval ← read(curr.val)) < v do
5:      prev ← curr
6:      curr ← read(curr.next)
7:    end while
8:    if tval = v then
9:      tnext ← read(curr.next)
10:     write(prev.next, tnext)
11:   Return (tval = v)
```

Algorithm 1.2 Using TM to implement remove operation of list-based set

```
1:  remove(v):
      atomic{
2:      prev ← head
3:      curr ← tx-read(prev.next)
4:      while (tval ← tx-read(curr.val)) < v do
5:        prev ← curr
6:        curr ← tx-read(curr.next)
7:      end while
8:      if tval = v then
9:        tnext ← tx-read(curr.next)
10:       tx-write(prev.next, tnext)
11:     tryCommit()
12:     Return (tval = v)
        } catch {AbortException a }
13:   { Return ⊥ }
```

Figure 1.1: Transforming a sequential implementation of the list-based set to a TM-based concurrent one

1.2 Transactional Memory (TM)

Transactional Memory (TM) [80, 117] addresses the challenge of resolving *conflicts* (concurrent reading and writing to the same data) in an efficient and safe manner by offering a simple interface in which sequences of shared memory operations on *data items* can be declared as *optimistic transactions*. The underlying idea of TM, inspired by databases [58], is to treat each transaction as *atomic*: a transaction may either *commit*, in which case it appears as executed sequentially, or *abort*, in which case none of its update operations *take effect*. Thus, it enables the programmer to design software applications having only sequential semantics in mind and let TM take care of *dynamically* handling the conflicts resulting from concurrent executions at run-time.

A TM *implementation* provides processes with algorithms for implementing *transactional operations* such as *read*, *write*, *tryCommit* and *tryAbort* on *data items* using base objects. TM implementations typically ensure that all committed transactions appear to execute sequentially in some total order respecting the timing of non-overlapping transactions. Moreover, unlike database transactions, intermediate states witnessed by the read operations of an incomplete transaction may affect the user application. Thus, to ensure that the TM implementation does not export any pathological executions, it is additionally expected that every transaction (including aborted and incomplete ones) must return responses that is consistent with some *sequential* execution of the TM implementation.

In general, given a sequential implementation of a data abstraction, a corresponding *TM-based* concurrent one encapsulates the sequential (high-level) operations within a transaction. Then, the TM-based concurrent implementation of the data abstraction replaces each read and write of a data item with the transactional read and write implementations, respectively. If the transaction commits, then the result of the high-level operation is returned to the application. Otherwise, one of the transactional operations may be aborted, in which case, the write operations performed by the transaction do not take effect and the high-level operation is typically re-started with a new transaction.

To illustrate this, we refer to the sequential implementation of the remove operation of list-based set depicted in Algorithm 1.1 of Figure 1.1. In a TM-based concurrent implementation of the list-based set (Algorithm 1.2), each read (and resp. write) operation performed by remove(v) on a data item X of the list is replaced with tx-read(X) (and resp., tx-write(X, arg)). tx-read(X) returns the value of the data item X or aborts the transaction while tx-write(X, arg) writes the value arg to X or aborts the transaction. Finally, the process attempts to commit the transaction by invoking the tryCommit operation. If the tryCommit is successful, the response of remove(v) is returned; otherwise a failed response (denoted $\perp$) is returned, in which case, the write operations performed by the transaction are "rolled back".

Intuitively, it is easy to understand how TM simplifies concurrent programming. Deriving a TM-based concurrent implementation of the list-based set simply requires encapsulating the operations to be executed atomically within a transaction using an *atomic* delimiter [1]. The underlying TM implementation endeavours to dynamically execute the transactions by resolving the conflicts that might arise from processes reading and writing to the same data item at run-time. Intuitively, since the TM implementation enforces a strong safety property, the resulting list-based implementation is also safe: the responses of its operations are consistent with some sequential execution of the list-based set.

One may view TM as a *universal construction* [13, 28, 45, 49, 75] that accepts as input the operations of a sequential implementation and strives to execute them concurrently. Specifically, TM is designed to work in a dynamic environment where neither the sequence of operations nor the data items accessed by a transaction are known a priori. Thus, the response of a read operation performed by a transaction is returned immediately to the application, and the application determines the next data item that must accessed by the transaction.

TM-based implementations overcome the drawbacks of traditional synchronization techniques based on locks and compare-and-swap. Firstly, the TM interface places minimal overhead on the programmer: using TM only requires encapsulating the sequential operations within transactions and handling an exception should the transaction be aborted. Secondly, the ability to execute multiple operations atomically allows TM-based implementations to seamlessly compose smaller atomic operations to produce larger ones. For example, suppose that we wish to atomically dequeue from a queue Q_1 and enqueue(v) in queue Q_2, where v is the value returned by Q_1.dequeue. Solving this problem using TM simply requires encapsulating the sequential implementation of Q_1.dequeue, followed by Q_2.enqueue(v) within a transaction.

Note that a TM implementation may internally employ locks or conditional rmw instructions like compare-and-swap. However, TM raises the level of abstraction by exposing an easy-to-use compositional transactional interface for the user application that is *oblivious* to the specifics of the implementation and the semantics of the user application.

1.3 Summary of the thesis

TM allows the programmer to *speculatively* execute sequences of shared-memory operations as atomic transactions: if the transaction commits, the operations appear as executed sequentially, or if the transaction aborts, the update operations do not take effect. The combination of speculation and the simple programming interface provided by TM seemingly overcomes the problems associated with traditional synchronization techniques based on locks and compare-and-swap. But are there some fundamental drawbacks associated with the TM abstraction? Does providing high degrees of concurrency in TMs come with inherent costs? This is the central question of the thesis. In the rest of this introductory chapter, we provide a summary of the results in the thesis that give some answers to this question.

1.3.1 Safety for transactional memory

We first need to define the consistency criteria that must be satisfied by a TM implementation. We formalize the semantics of a *safe* TM: every transaction, including aborted and incomplete ones, must observe a view that is *consistent* with some sequential execution. This is important, since if the intermediate view is not consistent with *any* sequential execution, the application may experience a fatal irrevocable error or enter an infinite loop. Additionally, the response of a transaction's read should not depend on an ongoing transaction that has not started committing yet. This restriction, referred to as *deferred-update semantics* appears desirable, since the ongoing transaction may still abort, thus rendering the read inconsistent. We define the notion of deferred-update semantics formally and apply it to several TM consistency criteria proposed in literature. We then verify if the resulting TM consistency

[1] Different compilers may use different names for the delimiter; in *GCC*, it is transaction_atomic [2]

criterion is a *safety* property [11, 100, 108] in the formal sense, *i.e.*, the set of *histories* (interleavings of invocations and responses of transactional operations) is *prefix-closed* and *limit-closed*.

We first consider the popular consistency criterion of *opacity* [64]. Opacity requires the states observed by all transactions, included uncommitted ones, to be consistent with a global *serialization*, *i.e.*, a serial execution constituted by committed transactions. Moreover, the serialization should respect the *real-time order*: a transaction that completed before (in real time) another transaction started should appear first in the serialization.

By definition, opacity reduces correctness of a TM history to correctness of all its prefixes, and thus is prefix-closed and limit-closed. Thus, to verify that a history is opaque, one needs to verify that each of its prefixes is consistent with some global serialization. To simplify verification and explicitly introduce deferred-update semantics into a TM correctness criterion, we specify a general criterion of *du-opacity* [18], which requires the global serial execution to respect the deferred-update property. Informally, a du-opaque history must be indistinguishable from a totally-ordered history, with respect to which no transaction reads from a transaction that has not started committing. Assuming that in an infinite history, every transaction completes each of the operations it invoked, we prove that du-opacity is a safety property.

One may notice that the intended safety semantics does not require, as opacity does, that all transactions observe the same serial execution. As long as committed transactions constitute a serial execution and every transaction witnesses a consistent state, the execution can be considered "safe": no run-time error that cannot occur in a serial execution can happen. Two definitions in literature have adopted this approach [43, 85]. We introduce "deferred-update" versions of these properties and discuss how the resulting properties relate to du-opacity.

1.3.2 Complexity of transactional memory

One may observe that a TM implementation that aborts or never commits any transaction is trivially safe, but not very useful. Thus, the TM implementation must satisfy some nontrivial *liveness* property specifying the conditions under which the transactional operations must return some response and a *progress* property specifying the conditions under which the transaction is allowed to abort.

Two properties considered important for TM performance are *read invisibility* [23] and *disjoint-access parallelism* [86]. Read invisibility may boost the concurrency of a TM implementation by ensuring that no reading transaction can cause any other transaction to abort. The idea of disjoint-access parallelism is to allow transactions that do not access the same data item to proceed independently of each other without memory contention.

We investigate the inherent complexities in terms of time and memory resources associated with implementing safe TMs that provide strong liveness and progress properties, possibly combined with attractive requirements like read invisibility and disjoint-access parallelism. Which classes of TM implementations are (im)possible to solve?

Blocking TMs. We begin by studying TM implementations that are *blocking*, in the sense that, a transaction may be delayed or aborted due to concurrent transactions.

- We prove that, even inherently *sequential* TMs, that allow a transaction to be aborted due to a concurrent transaction, incur significant complexity costs when combined with read invisibility and disjoint-access parallelism.
- We prove that, *progressive* TMs, that allow a transaction to be aborted only if it encounters a read-write or write-write conflict with a concurrent transaction [62], may need to exclusively control a linear number of data items at some point in the execution.
- We then turn our focus to *strongly progressive* TMs [64] that, in addition to progressiveness, ensures that *not all* concurrent transactions conflicting over a single data item abort. We prove that in any strongly progressive TM implementation that accesses the shared memory with *read*,

write and *conditional* primitives, such as compare-and-swap, the total number of *remote memory references* [14, 22] (RMRs) that take place in an execution in which n concurrent processes perform transactions on a single data item might reach $\Omega(n \log n)$ in the worst-case.

- We show that, with respect to the amount of *expensive synchronization* patterns like compare-and-swap instructions and *memory barriers* [17, 103], progressive implementations are asymptotically optimal. We use this result to establish a linear (in the transaction's data set size) separation between the worst-case transaction expensive synchronization complexity of progressive TMs and *permissive* TMs that allow a transaction to abort only if committing it would violate opacity.

Non-blocking TMs. Next, we focus on TMs that avoid using locks and rely on non-blocking synchronization: a prematurely halted transaction cannot not prevent other transactions from committing. Possibly the weakest non-blocking progress condition is obstruction-freedom [78, 82] stipulating that every transaction running in the absence of *step contention*, *i.e.*, not encountering steps of concurrent transactions, must commit. In fact, several early TM implementations [52, 79, 101, 117, 120] satisfied obstruction-freedom. However, circa. 2005, several papers presented the case for a shift from TMs that provide obstruction-free TM-progress to lock-based progressive TMs [39, 40, 48]. They argued that lock-based TMs tend to outperform obstruction-free ones by allowing for simpler algorithms with lower complexity overheads. We prove the following lower bounds for obstruction-free TMs.

- Combining invisible reads with even *weak* forms of disjoint-access parallelism [24] in obstruction-free TMs is impossible,
- A read operation in a n-process obstruction-free TM implementation incurs $\Omega(n)$ *memory stalls* [16, 46].
- A *read-only* transaction may need to perform a linear (in n) number of expensive synchronization patterns.

We then present a progressive TM implementation that beats all of these lower bounds, thus suggesting that the course correction from non-blocking (obstruction-free) TMs to blocking (progressive) TMs was indeed justified.

Partially non-blocking TMs. Lastly, we explore the costs of providing non-blocking progress to only a *subset* of transactions. Specifically, we require *read-only* transactions to commit *wait-free*, *i.e.*, every transaction commits within a finite number of its steps, but *updating* transactions are guaranteed to commit only if they run in the absence of concurrency. We show that combining this kind of *partial* wait-freedom with read invisibility *or* disjoint-access parallelism comes with inherent costs. Specifically, we establish the following lower bounds for TMs that provide this kind of partial wait-freedom.

- This kind of partial wait-freedom equipped with invisible reads results in maintaining unbounded sets of *versions* for every data item.
- It is impossible to implement a *strict* form of disjoint-access parallelism [60].
- Combining with the weak form of disjoint-access parallelism means that a read-only transaction (with an arbitrarily large read set) must sometimes perform at least one expensive synchronization pattern per read operation in some executions.

1.3.3 Hybrid transactional memory

We turn our focus on *Hybrid transactional memory (HyTM)* [35, 37, 88, 99]. The TM abstraction, in its original manifestation, augmented the processor's *cache-coherence protocol* and extended the CPU's instruction set with instructions to indicate which memory accesses must be transactional [80]. Most popular TM designs, subsequent to the original proposal in [80] have implemented all the functionality in software [36, 52, 79, 101, 117]. More recently, CPUs have included hardware extensions to support small transactions [1, 107, 111]. Hardware transactions may be spuriously aborted due to several reasons: cache capacity overflow, interrupts *etc.* This has led to proposals for *best-effort* HyTMs in which the fast, but potentially unreliable hardware transactions are complemented with slower, but more reliable software

transactions. However, the fundamental limitations of building a HyTM with nontrivial concurrency between hardware and software transactions are not well understood. Typically, hardware transactions usually employ *code instrumentation* techniques to detect concurrency scenarios and abort in the case of contention. But are there inherent instrumentation costs of implementing a HyTM, and what are the trade-offs between these costs ands provided concurrency, *i.e.*, the ability of the HyTM to execute hardware and software transactions in parallel?

The thesis makes the following contributions which help determine the cost of concurrency in HyTMs.

- We propose a general model for HyTM implementations, which captures the notion of *cached* accesses as performed by hardware transactions, and precisely defines instrumentation costs in a quantifiable way.
- We derive lower and upper bounds in this model, which capture for the first time, an inherent trade-off on the degree of concurrency allowed between hardware and software transactions and the *instrumentation* overhead introduced on the hardware.

1.3.4 Optimism for boosting concurrency

Lock-based implementations are conventionally *pessimistic* in nature: the operations invoked by processes are not "abortable" and return only after they are successfully completed. The TM abstraction is a realization of *optimistic* concurrency control: speculatively execute transactions, abort and roll back on dynamically detected conflicts. But are optimistic implementations fundamentally better equipped to exploit concurrency than pessimistic ones?

We compare the *amount of concurrency* one can obtain by converting a sequential implementation of a data abstraction into a concurrent one using optimistic or pessimistic synchronization techniques. To establish fair comparison of such implementations, we introduce a new correctness criterion for concurrent implementations, called *locally serializable linearizability*, defined independently of the synchronization techniques they use.

We treat an implementation's concurrency as its ability to accept *schedules* of sequential operations from different processes. More specifically, we assume an external scheduler that defines which processes execute which steps of the corresponding sequential implementation in a dynamic and unpredictable fashion. This allows us to define concurrency provided by an implementation as the set of interleavings of steps of sequential operations (or schedules) it *accepts*, *i.e.*, is able to effectively process. Then, the more schedules the implementation would accept without hampering correctness, the more concurrent it would be.

The thesis makes the following contributions.

- We provide a framework to analytically capture the inherent concurrency provided by two broad classes of synchronization techniques: pessimistic implementations that implement some form of mutual exclusion and optimistic implementations based on speculative executions.
- We show that, implementations based on pessimistic synchronization and "semantics-oblivious" TMs are suboptimal, in the sense that, there exist there exist simple schedules of the list-based set which cannot be accepted by *any* pessimistic or TM-based implementation. Specifically, we prove that TM-based implementations accept schedules of the list-based set that cannot be accepted by *any* pessimistic implementation. However, we also show pessimistic implementations of the list-based set which accept schedules that cannot be accepted by *any* TM-based implementation.
- We show that, there exists an optimistic implementation of the list-based set that is *concurrency optimal*, *i.e.*, it accepts *all* correct schedules.

Our results suggest that "semantics-aware" optimistic implementations may be better suited to exploiting concurrency than their pessimistic counterparts.

1.4 Roadmap of the thesis

We first define the TM model, the TM properties proposed in literature and the complexity metrics considered in Chapter 2. Chapter 3 is on safety for TMs. Chapter 4 is on the complexity of blocking TMs, non-blocking TMs that satisfy obstruction-freedom are covered in Chapter 5 and we present lower bounds for partially non-blocking TMs in Chapter 6. Chapter 7 is devoted to the study of hybrid TMs. In Chapter 8, we compare the relative abilities of optimistic and pessimistic synchronization techniques in exploiting concurrency. Chapter 9 presents closing comments and future directions. Viewed collectively, the results hopefully shine light on the foundations of the TM abstraction that is widely expected to be the *Zeitgeist* of the concurrent computational model.

2

Transactional Memory model

> All models are wrong, but some models are useful.
>
> *George Edward Pelham Box*

In this chapter, we formalize the TM model and discuss some important TM properties proposed in literature. In Section 2.1, we formalize the specification of TMs. In Section 2.2, we introduce the basic TM-correctness property of *strict serializability* that we consider in the thesis. Sections 2.3 and 2.4 overview progress and liveness properties for TMs respectively and identifies the relations between them. Section 2.5 defines the notion of *invisible reads* while Section 2.6 is on *disjoint-access parallelism*. Finally, in Section 2.7, we introduce some of the complexity metrics considered in the thesis.

2.1 TM interface and TM implementations

In this section, we first describe the *shared memory* model of computation and then introduce the TM abstraction.

The shared memory model. The thesis considers the standard *asynchronous shared memory* model of computation in which a set of $n \in \mathbb{N}$ processes (that may *fail by crashing*), communicate by applying *operations* on shared *objects* [16]. An object is an instance of an *abstract data type*. An abstract data type τ is a *mealy machine* that is specified as a tuple $(\Phi, \Gamma, Q, q_0, \delta)$ where Φ is a set of operations, Γ is a set of responses, Q is a set of states, $q_0 \in Q$ is an initial state and $\delta \subseteq Q \times \Phi \times Q \times \Gamma$ is a transition relation that determines, for each state and each operation, the set of possible resulting states and produced responses [6]. Here, $(q, \pi, q', r) \in \delta$ implies that when an operation $\pi \in \Phi$ is applied on an object of type τ in state q, the object moves to state q' and returns response r.

An *implementation* of an object type τ provides a specific data-representation of τ that is realized by processes applying *primitives* on shared *base objects*, each of which is assigned an initial value. In order to implement an object, processes are provided with an algorithm, which is a set of deterministic state-machines, one for each process. In the thesis, we use the term primitive to refer to operations on base objects and reserve the term operation for the object that is implemented from the base objects.

A primitive is a generic atomic *read-modify-write* (*rmw*) procedure applied to a base object [46, 75]. It is characterized by a pair of functions $\langle g, h \rangle$: given the current state of the base object, g is an *update function* that computes its state after the primitive is applied, while h is a *response function* that specifies the outcome of the primitive returned to the process. A rmw primitive is *trivial* if it never changes the value of the base object to which it is applied. Otherwise, it is *nontrivial*. An rmw primitive $\langle g, h \rangle$ is *conditional* if there exists v, w such that $g(v, w) = v$ and there exists v, w such that $g(v, w) \neq v$ [51].

Read is an example of a trivial rmw primitive that takes no input arguments: when applied to a base object with value v, the update function leaves the state of the base object unchanged and the response function returns the value v. *Write* is an example of a nontrivial rmw primitive that takes an input argument v': when applied to a base object with value v, its update function changes the value of the base object to v' and its response function returns *ok*. *Compare-and-swap* is an example of a nontrivial conditional rmw primitive: its update function receives an input argument $\langle old, new \rangle$ and changes the value v of the base object to which it is applied *iff* $v = old$. *Load-linked/store-conditional* is another example of a nontrivial conditional rmw primitive: the *load-linked* primitive executed by some process p_i returns the value of the base object to which it is applied and the *store-conditional* primitive's update function receives an input *new* and atomically changes the value of the base object to *new iff* the base object has not been updated by any other process since the load-linked event by p_i. *Fetch-and-add* is an example of a nontrivial rmw primitive that is not conditional: its update function applied to base object with an integer value v takes an integer w as input and changes the value of the base object to $v + w$.

Transactional memory (TM). *Transactional memory* allows a set of data items (called *t-objects*) to be accessed via *transactions*. Every transaction T_k has a unique identifier k. We make no assumptions on the *size* of a t-object, *i.e.*, the cardinality on the set V of possible values a t-object can store. A transaction T_k may contain the following *t-operations*, each being a matching pair of an *invocation* and a *response*: $read_k(X)$ returns a value in V, denoted $read_k(X) \to v$, or a special value $A_k \notin V$ (*abort*); $write_k(X, v)$, for a value $v \in V$, returns *ok* or A_k; $tryC_k$ returns $C_k \notin V$ (*commit*) or A_k. As we show in the subsequent Section 2.2, we can specify TM as an abstract data type.

Note that a TM interface may additionally provide a $start_k$ t-operation that returns *ok* or A_k, which is the first t-operation transaction T_k must invoke, or a $tryA_k$ t-operation that returns A_k. However, the actions performed inside the $start_k$ may be performed as part of the first t-operation performed by the transaction. The $tryA_k$ t-operation allows the user application to explicitly abort a transaction and can be useful, but since each of the individual t-read or t-write are allowed to abort, the $tryA_k$ t-operation provides no additional expressive power to the TM interface. Thus, for simplicity, we do not incorporate these t-operations in our TM specification.

TM implementations. A TM *implementation* provides processes with algorithms for implementing $read_k$, $write_k$ and $tryC_k()$ of a transaction T_k by applying primitives from a set of shared base objects, each of which is assigned an initial value. We assume that a process starts a new transaction only after its previous transaction has committed or aborted.

In the rest of this section, we define the terms specifically in the context of TM implementations, but they may be used analogously in the context of any concurrent implementation of an abstract data type.

Executions and configurations. An *event* of a process p_i (sometimes we say *step* of p_i) is an invocation or response of an operation performed by p_i or a rmw primitive $\langle g, h \rangle$ applied by p_i to a base object b along with its response r (we call it a *rmw event* and write $(b, \langle g, h \rangle, r, i)$).

A *configuration* (of an implementation) specifies the value of each base object and the state of each process. The *initial configuration* is the configuration in which all base objects have their initial values and all processes are in their initial states.

An *execution fragment* is a (finite or infinite) sequence of events. An *execution* of an implementation M is an execution fragment where, starting from the initial configuration, each event is issued according to M and each response of a rmw event $(b, \langle g, h \rangle, r, i)$ matches the state of b resulting from all preceding events. An execution $E \cdot E'$ denotes the concatenation of E and execution fragment E', and we say that E' is an *extension* of E or E' *extends* E.

Let E be an execution fragment. For every transaction (resp., process) identifier k, $E|k$ denotes the subsequence of E restricted to events of transaction T_k (resp., process p_k). If $E|k$ is non-empty, we say that T_k (resp., p_k) *participates* in E, else we say E is T_k*-free* (resp., p_k*-free*). Two executions E and E' are *indistinguishable* to a set $\mathcal{T}$ of transactions, if for each transaction $T_k \in \mathcal{T}$, $E|k = E'|k$. A TM *history* is the subsequence of an execution consisting of the invocation and response events of t-operations. Two histories H and H' are *equivalent* if $\mathit{txns}(H) = \mathit{txns}(H')$ and for every transaction $T_k \in \mathit{txns}(H)$, $H|k = H'|k$.

Data sets of transactions. The *read set* (resp., the *write set*) of a transaction T_k in an execution E, denoted $\mathit{Rset}_E(T_k)$ (resp., $\mathit{Wset}_E(T_k)$), is the set of t-objects that T_k reads (resp., writes to) in E. More specifically, if E contains an invocation of $\mathit{read}_k(X)$ (resp., $\mathit{write}_k(X, v)$), we say that $X \in \mathit{Rset}_E(T_k)$ (resp., $\mathit{Wset}_E(T_k)$) (for brevity, we sometimes omit the subscript E from the notation). The *data set* of T_k is $\mathit{Dset}(T_k) = \mathit{Rset}(T_k) \cup \mathit{Wset}(T_k)$. A transaction is called *read-only* if $\mathit{Rset}(T_k) \neq \emptyset \wedge \mathit{Wset}(T_k) = \emptyset$; *write-only* if $\mathit{Wset}(T_k) \neq \emptyset \wedge \mathit{Rset}(T_k) = \emptyset$ and *updating* if $\mathit{Wset}(T_k) \neq \emptyset$. Note that, in our TM model, the data set of a transaction is not known apriori, *i.e.*, at the start of the transaction and it is identifiable only by the set of data items the transaction has invoked a read or write on in the given execution.

Transaction orders. Let $\mathit{txns}(E)$ denote the set of transactions that participate in E. In an infinite history H, we assume that each $T_k \in \mathit{txns}(H)$, $H|k$ is finite; *i.e.*, transactions do not issue an infinite number of t-operations. An execution E is *sequential* if every invocation of a t-operation is either the last event in the history H exported by E or is immediately followed by a matching response. We assume that executions are *well-formed*, *i.e.*, for all T_k, $E|k$ begins with the invocation of a t-operation, is sequential and has no events after A_k or C_k. A transaction $T_k \in \mathit{txns}(E)$ is *complete in* E if $E|k$ ends with a response event. The execution E is *complete* if all transactions in $\mathit{txns}(E)$ are complete in E. A transaction $T_k \in \mathit{txns}(E)$ is *t-complete* if $E|k$ ends with A_k or C_k; otherwise, T_k is *t-incomplete*. T_k is *committed* (resp., *aborted*) in E if the last event of T_k is C_k (resp., A_k). The execution E is *t-complete* if all transactions in $\mathit{txns}(E)$ are t-complete.

For transactions $\{T_k, T_m\} \in \mathit{txns}(E)$, we say that T_k *precedes* T_m in the *real-time order* of E, denoted $T_k \prec_E^{RT} T_m$, if T_k is t-complete in E and the last event of T_k precedes the first event of T_m in E. If neither $T_k \prec_E^{RT} T_m$ nor $T_m \prec_E^{RT} T_k$, then T_k and T_m are *concurrent* in E. An execution E is *t-sequential* if there are no concurrent transactions in E.

Latest written value and legality. Let H be a t-sequential history. For every operation $\mathit{read}_k(X)$ in H, we define the *latest written value* of X as follows: if T_k contains a $\mathit{write}_k(X, v)$ preceding $\mathit{read}_k(X)$, then the latest written value of X is the value of the latest such write to X. Otherwise, the latest written value of X is the value of the argument of the latest $\mathit{write}_m(X, v)$ that precedes $\mathit{read}_k(X)$ and belongs to a committed transaction in H. (This write is well-defined since H starts with T_0 writing to all t-objects.)

We say that $\mathit{read}_k(X)$ is *legal* in a t-sequential history H if it returns the latest written value of X, and H is *legal* if every $\mathit{read}_k(X)$ in H that does not return A_k is legal in H.

We also assume, for simplicity, that the user application invokes a $\mathit{read}_k(X)$ at most once within a transaction T_k. This assumption incurs no loss of generality, since a repeated read can be assigned to return a previously returned value without affecting the history's legality.

Contention. We say that a configuration C after an execution E is *quiescent* (resp., *t-quiescent*) if every transaction $T_k \in \mathit{txns}(E)$ is complete (resp., t-complete) in C. If a transaction T is incomplete in an execution E, it has exactly one *enabled* event, which is the next event the transaction will perform according to the TM implementation. Events e and e' of an execution E *contend* on a base object b if they are both events on b in E and at least one of them is nontrivial (the event is trivial (resp., nontrivial) if it is the application of a trivial (resp., nontrivial) primitive).

We say that T is *poised to apply an event* e *after* E if e is the next enabled event for T in E. We say that transactions T and T' *concurrently contend on* b *in* E if they are poised to apply contending events on b after E.

We say that an execution fragment E is *step contention-free for t-operation* op_k if the events of $E|op_k$ are contiguous in E. We say that an execution fragment E is *step contention-free for* T_k if the events of $E|k$ are contiguous in E. We say that E is *step contention-free* if E is step contention-free for all transactions that participate in E.

2.2 TM-correctness

Correctness for TMs is specified as a safety property on TM histories [11, 100, 108]. In this section, we introduce the popular TM-correctness condition *strict serializability* [109]: all committed transactions appear to execute sequentially in some total order respecting the real-time transaction orders. We then explain how strict serializability is related to *linearizability* [83].

In the thesis, we only consider TM-correctness conditions like strict serializability and its restrictions. We formally define strict serializability below, but other TM-correctness conditions studied in the thesis can be found in Chapter 3.

First, we define how to derive a t-complete history from a t-incomplete one.

Definition 2.1 (Completions)**.** *Let H be a history. A* completion of H*, denoted $\overline{H}$, is a history derived from H as follows:*

- *First, for every transaction $T_k \in txns(H)$ with an incomplete t-operation op_k in H, if $op_k = read_k \vee write_k$, insert A_k somewhere after the invocation of op_k; otherwise, if $op_k = tryC_k()$, insert C_k or A_k somewhere after the last event of T_k.*
- *After all transactions are complete, for every transaction T_k that is not t-complete, insert $tryC_k \cdot A_k$ after the last event of transaction T_k.*

Definition 2.2 (Strict serializability)**.** *A finite history H is* serializable *if there is a legal t-complete t-sequential history S, such that there is a completion $\overline{H}$ of H, such that S is equivalent to $cseq(\overline{H})$, where $cseq(\overline{H})$ is the subsequence of $\overline{H}$ reduced to committed transactions in $\overline{H}$.*

We refer to S as a serialization *of H.*

We say that H is strictly serializable *if there exists a serialization S of H such that for any two transactions $T_k, T_m \in txns(H)$, if $T_k \prec_H^{RT} T_m$, then T_k precedes T_m in S.*

In general, given a TM-correctness condition C, we say that a TM implementation M satisfies C if every execution of M satisfies C.

Strict serializability as linearizability. We now show we can specify TM as an abstract data type.

The sequential specification of a TM is specified as follows:

1. Φ is the set of all transactions $\{T_i\}_{i \in \mathbb{N}}$
2. Γ is the set of incommensurate vectors $\{[r_1, \ldots, r_i]\}; i \in \mathbb{N}$; where each r_j; $1 \leq j \leq i-1 \in \{v \in V, A, ok\}$ and $r_i \in \{A, C\}$
3. The state of TM is a vector of the state of each t-object X_m. The state of a t-object X_m is a value $v_m \in V$ of X_m. Thus, $Q \subseteq \{[v_1^i, \ldots, v_m^i, \ldots]\}$; where each $v_m^i \in V$
4. $q_0 \in Q = [ov_1, \ldots, ov_m, \ldots]$, where each $ov_m \in V$
5. δ is defined as follows: Let T_k be a transaction applied to the TM in state $q = [v_1, \ldots, v_m, \ldots]$.
   - For every $X \in Rset(T_k)$, the response of $read_k(X)$ is defined as follows: If T_k contains a $write_k(X, v)$ prior to $read_k(X)$, then the response is v; else the response is the current state of X.
   - For every $X \in Wset(T_k)$, the response of $write_k(X, v)$ is ok.

- Transaction T_k returns the response C in which case the TM moves to state q' defined as follows: every $X_j \in Wset(T_k)$ to which T_k writes values nv_j, $q'[j] = nv_j$; else if $X_j \notin Wset(T_k)$, $q'[j]$ is unchanged. Otherwise, T_k returns the response A in which case $q' = q$.

In general, the correctness of an implementation of a data type is commonly captured by the criterion of linearizability. In the TM context, a t-complete history H is *linearizable* with respect to the TM type if there exists a t-sequential history S equivalent to H such that (1) S respects the real-time ordering of transactions in H and (2) *S is consistent with the sequential specification of TM.*

The following lemma, which illustrates the similarity between strict serializability and linearizability with respect to the TM type, is now immediate.

Lemma 2.1. *Let H be any t-complete history. Then, H is strictly serializable* iff *H is linearizable with respect to the TM type.*

2.3 TM-progress

One may notice that a TM implementation that forces, in every execution to abort every transaction is trivially strictly serializable, but not very useful. A TM-progress condition specifies the conditions under which a transaction is allowed to abort. Technically, a TM-progress condition specified this way is a *safety property* since it can be violated in a finite execution (cf. Chapter 3 for details on safety properties).

Ideally, a TM-progress condition must provide *non-blocking* progress, in the sense that a prematurely halted transaction cannot prevent all other transactions from committing. Such a TM-progress condition is also said to be *lock-free* since it cannot be achieved by use of locks and mutual-exclusion. A non-blocking TM-progress condition is considered useful in asynchronous systems with process failures since it prevents the TM implementation from deadlocking (processes wait infinitely long without committing their transactions).

Obstruction-freedom. Perhaps, the weakest non-blocking TM-progress condition is obstruction-freedom, which stipulates that a transaction may be aborted only if it encounters steps of a concurrent transaction [82].

Definition 2.3 (Obstruction-free (OF) TM-progress)**.** *We say that a TM implementation M provides* obstruction-free (OF) TM-progress *if for every execution E of M, if any transaction $T_k \in txns(E)$ returns A_k in E, then E is not step contention-free for T_k.*

We now survey the popular *blocking* TM-progress properties proposed in literature. Intuitively, unlike non-blocking TM-progress conditions that adapt to *step contention*, a blocking TM-progress condition allows a transaction to be aborted due to *overlap contention.*

Minimal progressiveness. Intuitively, the most basic TM-progress condition is one which provide only *sequential* TM-progress, *i.e.*, a transaction may be aborted due to a concurrent transaction. In literature, this is referred to as *minimal progressiveness* [64].

Definition 2.4 (Minimal progressiveness)**.** *We say that a TM implementation M provides* minimal progressive *TM-progress (or* minimal progressiveness*) if for every execution E of M and every transaction $T_k \in txns(E)$ that returns A_k in E, there exists a transaction $T_m \in txns(E)$ that is concurrent to T_k in E [64].*

Given TM conditions C_1 and C_2, if every TM implementation that satisfies C_1 also satisfies C_2, but the converse is not true, we say that $C_2 \ll C_1$.

Observation 2.2. *Minimal progressiveness $\ll$ Obstruction-free.*

Proof. Clearly, every TM implementation that satisfies obstruction-freedom also satisfies minimal progressiveness, but the converse is not true. Consider any execution of a TM implementation M in which a transaction T run step contention-free. If M is minimally progressive, then T may be aborted in such an execution since T may be concurrent with another transaction. However, if M satisfies obstruction-freedom, T cannot be aborted in such an execution. □

Progressiveness. In contrast to the "single-lock" minimal progressive TM-progress condition (also referred to as *sequential* TM-progress in the thesis), state-of-the-art TM implementations allow a transaction to abort only if it encounters a *conflict* on a t-object with a concurrent transaction.

Definition 2.5 (Conflicts)**.** *We say that transactions T_i, T_j* conflict *in an execution E on a t-object X if T_i and T_j are concurrent in E and $X \in Dset(T_i) \cap Dset(T_j)$, and $X \in Wset(T_i) \cup Wset(T_j)$.*

Definition 2.6 (Progressiveness)**.** *A TM implementation M provides* progressive *TM-progress (or* progressiveness*) if for every execution E of M and every transaction $T_i \in txns(E)$ that returns A_i in E, there exists a transaction $T_k \in txns(E)$ such that T_k and T_i conflict in E [64].*

Note that progressiveness is incomparable to obstruction-freedom.

Observation 2.3. *Progressiveness $\not\ll$ Obstruction-free and Obstruction-free $\not\ll$ Progressiveness.*

Proof. We can show that there exists an execution exported by an obstruction-free TM, but not by any progressive TM and vice-versa.

Consider a t-read X by a transaction T that runs step contention-free from a configuration that contains an incomplete write to X. Weak progressiveness does not preclude T from being aborted in such an execution. Obstruction-free TMs however, must ensure that T must complete its read of X without blocking or aborting in such executions. On the other hand, weak progressiveness requires two non-conflicting transactions to not be aborted even in executions that are not step contention-free; but this is not guaranteed by obstruction-freedom. □

In general, progressive TMs (including the ones described in the thesis) satisfy the following stronger definition: for every transaction $T_i \in txns(E)$ that returns A_i in an execution E, there exists prefix E' of E and a transaction $T_k \in txns(E')$ such that T_k and T_i conflict in E. However, for the lower bound results stated in the thesis, we stick to Definition 2.6.

Strong progressiveness. One may observe that the definition of progressiveness does not preclude two conflicting transactions (over a single t-object) from each being aborted. Thus, we study a stronger notion of progressiveness called *strong progressiveness* [64].

Let $CObj_E(T_i)$ denote the set of t-objects over which transaction $T_i \in txns(H)$ conflicts with any other transaction in an execution E, *i.e.*, $X \in CObj_E(T_i)$, *iff* there exist transactions T_i and T_k that conflict on X in E. Let $Q \subseteq txns(E)$ and $CObj_E(Q) = \bigcup_{T_i \in Q} CObj_E(T_i)$.

Let $CTrans(E)$ denote the set of non-empty subsets of $txns(E)$ such that a set Q is in $CTrans(E)$ if no transaction in Q conflicts with a transaction not in Q.

Definition 2.7 (Strong progressiveness)**.** *A TM implementation M is* strongly progressive *if M is weakly progressive and for every execution E of M and for every set $Q \in CTrans(E)$ such that $|CObj_E(Q)| \leq 1$, some transaction in Q is not aborted in E [64].*

The above definitions imply:

Corollary 2.4. *Minimal progressiveness (sequential TM-progress) $\ll$ Progressiveness $\ll$ Strong progressiveness.*

Mv-permissiveness. Perelman *et al.* introduced the notion of *mv-permissiveness*, a TM-progress property designed to prevent read-only transactions from being aborted.

Definition 2.8 (Mv-permissiveness)**.** *A TM implementation M is* mv-permissive *if for every execution E of M and for every transaction $T_k \in txns(E)$ that returns A_k in E, we have that $Wset(T_k) \neq \emptyset$ and there exists an updating transaction $T_m \in txns(E)$ such that T_k and T_m conflict in E.*

We observe that mv-permissiveness is strictly stronger than progressiveness, but incomparable to strong progressiveness.

Observation 2.5. *Progressiveness $\ll$ Mv-permissiveness.*

Proof. Since mv-permissive TMs allow a transaction to be aborted only on read-write conflicts, they also satisfy progressiveness. But the converse is not true. Consider an execution in which a read-only transaction T_i that runs concurrently with a conflicting updating transaction T_j. By the definition of progressiveness, both T_i and T_j may be aborted in such an execution. However, a mv-permissive TM would not allow T_i to be aborted since it is read-only. □

Observation 2.6. *Strong progressiveness $\not\ll$ Mv-permissiveness and Mv-permissiveness $\not\ll$ Strong progressiveness.*

Proof. Consider an execution in which a read-only transaction T_i that runs concurrently with an updating transaction T_j such that T_i and T_j conflict on at least two t-objects. By the definition of strong progressiveness, both T_i and T_j may be aborted in such an execution. However, a mv-permissive TM would not allow T_i to be aborted since it is read-only.

On the other hand, consider an execution in which two updating transactions T_i and T_j that conflict on a single t-object. A mv-permissive TM allows both T_i and T_j to be aborted, but strong progressiveness ensures that at least one of T_i or T_j is not aborted in such an execution. □

2.4 TM-liveness

Observe that a TM-progress condition only specifies the conditions under which a transaction is aborted, but does not specify the conditions under which it must commit. For instance, the OF TM-progress condition specifies that a transaction T may be aborted only in executions that are not step contention-free for T, but does not guarantee that T is committed in a step contention-free execution. Thus, in addition to a progress condition, we must stipulate a *liveness* [11, 100] condition.

We now define the TM-liveness conditions considered in the thesis.

Definition 2.9 (Sequential TM-liveness)**.** *A TM implementation M provides* sequential TM-liveness *if for every finite execution E of M, and every transaction T_k that runs t-sequentially and applies the invocation of a t-operation op_k immediately after E, the finite step contention-free extension for op_k contains a response.*

Definition 2.10 (Interval contention-free (ICF) TM-liveness)**.** *A TM implementation M provides* interval contention-free (ICF) *TM-liveness if for every finite execution E of M such that the configuration after E is quiescent, and every transaction T_k that applies the invocation of a t-operation op_k immediately after E, the finite step contention-free extension for op_k contains a response.*

Definition 2.11 (Starvation-free TM-liveness)**.** *A TM implementation M provides* starvation-free TM-liveness *if in every execution of M, each t-operation eventually returns a matching response, assuming that no concurrent t-operation stops indefinitely before returning.*

Definition 2.12 (Obstruction-free (OF) TM-liveness)**.** *A TM implementation M provides* obstruction-free (OF) TM-liveness *if for every finite execution E of M, and every transaction T_k that applies the invocation of a t-operation op_k immediately after E, the finite step contention-free extension for op_k contains a matching response.*

Definition 2.13 (Wait-free (WF) TM-liveness)**.** *A TM implementation M provides* wait-free (WF) TM-liveness *if in every execution of M, every t-operation returns a response in a finite number of its steps.*

The following observations are immediate from the definitions:

Observation 2.7. *Sequential TM-liveness $\ll$ ICF TM-liveness $\ll$ OF TM-liveness $\ll$ WF TM-liveness and Starvation-free TM-liveness $\ll$ WF TM-liveness.*

Since ICF TM-liveness guarantees that a t-operation returns a response if there is no other concurrent t-operation, we have:

Observation 2.8. *ICF TM-liveness $\ll$ Starvation-free TM-liveness.*

However, we observe that OF TM-liveness and starvation-free TM-liveness are incomparable.

Observation 2.9. *Starvation-free TM-liveness $\not\ll$ OF TM-liveness and OF TM-liveness $\not\ll$ Starvation-free TM-liveness.*

Proof. Consider the step contention-free execution of t-operation op_k concurrent with t-operation op_m: op_k must return a matching response within a finite number of its steps, but this is not necessarily ensured by starvation-free TM-liveness (op_m may be delayed indefinitely). On the other hand, in executions where two concurrent t-operations op_k and op_k encounter step contention, but neither stalls indefinitely, both must return matching responses. But this is not guaranteed by OF TM-liveness. □

2.5 Invisible reads

In this section, we introduce the notion of *invisible reads* that intuitively ensures that a reading transaction does not cause a concurrent transaction to abort. Since most TM worklods are believed to be read-dominated, this is considered to be an important TM property for performance [25, 65].

Invisible reads. Informally, in a TM using invisible reads, a transaction cannot reveal any information about its read set to other transactions. Thus, given an execution E and some transaction T_k with a non-empty read set, transactions other than T_k cannot distinguish E from an execution in which T_k's read set is empty. This prevents TMs from applying nontrivial primitives during t-read operations and from announcing read sets of transactions during tryCommit. Most popular TM implementations like *TL2* [39] and *NOrec* [36] satisfy this property.

Definition 2.14 (Invisible reads [23])**.** *We say that a TM implementation M uses* invisible reads *if for every execution E of M:*

- *for every read-only transaction $T_k \in txns(E)$, no event of $E|k$ is nontrivial in E,*
- *for every updating transaction $T_k \in txns(E)$; $Rset_E(T_k) \neq \emptyset$, there exists an execution E' of M such that*
  - *$Rset_{E'}(T_k) = \emptyset$,*
  - *$txns(E) = txns(E')$ and $\forall T_m \in txns(E) \setminus \{T_k\}$: $E|m = E'|m$*
  - *for any two transactions $T_i, T_j \in txns(E)$, if the last event of T_i precedes the first event of T_j in E, then the last event of T_i precedes the first event of T_j in E'.*

Weak invisible reads. We introduce the notion of *weak* invisible reads that prevents t-read operations from applying nontrivial primitives only in the absence of concurrent transactions. Specifically, weak read invisibility allows t-read operations of a transaction T to be "visible", *i.e.*, write to base objects, only if T is concurrent with another transaction.

Definition 2.15 (Weak invisible reads)**.** *For any execution E and any t-operation π_k invoked by some transaction $T_k \in txns(E)$, let $E|\pi_k$ denote the subsequence of E restricted to events of π_k in E.*

We say that a TM implementation M satisfies weak invisible reads *if for any execution E of M and every transaction $T_k \in txns(E)$; $Rset(T_k) \neq \emptyset$ that is not concurrent with any transaction $T_m \in txns(E)$, $E|\pi_k$ does not contain any nontrivial events, where π_k is any t-read operation invoked by T_k in E.*

For example, the popular TM implementation *DSTM* [79] satisfies weak invisible reads, but not invisible reads. Algorithm 5.1 in Chapter 4 depicts a TM implementation that is based on DSTM satisfying weak invisible reads, but not the stronger definition of invisible reads.

2.6 Disjoint-access parallelism (DAP)

The notion of *disjoint-access parallelism (DAP)* [86] is considered important in the TM context since it allows two transactions accessing unrelated t-objects to execute without memory contention. In this section, we preview the DAP definitions proposed in literature and identify the relations between them.

Strict data-partitioning. Let $E|X$ denote the subsequence of the execution E derived by removing all events associated with t-object X. A TM implementation M is *strict data-partitioned* [64], if for every t-object X, there exists a set of base objects $Base_M(X)$ such that

- for any two t-objects X_1, X_2; $Base_M(X_1) \cap Base_M(X_2) = \emptyset$,
- for every execution E of M and every transaction $T \in txns(E)$, every base object accessed by T in E is contained in $Base_M(X)$ for some $X \in Dset(T)$
- for all executions E and E' of M, if $E|X = E|X$ for some t-object X, then the configurations after E and E' only differ in the states of the base objects in $Base_M(X)$.

Strict disjoint-access parallelism. A TM implementation M is *strictly disjoint-access parallel (strict DAP)* if, for all executions E of M, and for all transactions T_i and T_j that participate in E, T_i and T_j contend on a base object in E only if $Dset(T_i) \cap Dset(T_j) \neq \emptyset$ [64].

Proposition 2.1. *Strict DAP $\ll$ Strict data-partitioning.*

Proof. Let M be any strict data-partitioned TM implementation. Then, M is also strict DAP. Indeed, since any two transactions accessing mutually disjoint data sets in a strict data-partitioned implementation cannot access a common base object in any execution E of M, E also ensures that any two transactions contend on the same base object in E only if they access a common t-object.

Consider the following execution E of a strict DAP TM implementaton M that begins with two transactions T_1 and T_2 that access disjoint data sets in E. A strict data-partitioned TM implementation would preclude transactions T_1 and T_2 from accessing the same base object, but a strict DAP TM implementation does not preclude this possibility. □

We now describe two relaxations of strict DAP. For the formal definitions, we introduce the notion of a *conflict graph* which captures the dependency relation among t-objects accessed by transactions.

Read-write (RW) disjoint-access parallelism. Informally, read-write (RW) DAP means that two transactions can *contend* on a common base object only if their data sets are connected in the *conflict graph*, capturing write-set overlaps among all concurrent transactions.

We denote by $\tau_E(T_i, T_j)$, the set of transactions (T_i and T_j included) that are concurrent to at least one of T_i and T_j in an execution E.

Let $\tilde{G}(T_i, T_j, E)$ be an undirected graph whose vertex set is $\bigcup_{T \in \tau_E(T_i, T_j)} Dset(T)$ and there is an edge between t-objects X and Y *iff* there exists $T \in \tau_E(T_i, T_j)$ such that $\{X, Y\} \in Wset(T)$. We say that

T_i and T_j are *read-write disjoint-access* in E if there is no path between a t-object in $Dset(T_i)$ and a t-object in $Dset(T_j)$ in $\tilde{G}(T_i, T_j, E)$. A TM implementation M is *read-write disjoint-access parallel (RW DAP)* if, for all executions E of M, transactions T_i and T_j contend on the same base object in E only if T_i and T_j are not read-write disjoint-access in E or there exists a t-object $X \in Dset(T_i) \cap Dset(T_j)$.

Proposition 2.2. *RW DAP $\ll$ Strict DAP.*

Proof. From the definitions, it is immediate that every strict DAP TM implementation satisfies RW DAP.

But the converse is not true (Algorithm 5.1 describes a TM implementation that satisfies RW and weak DAP, but not strict DAP). Consider the following execution E of a weak DAP or RW DAP TM implementaton M that begins with the t-incomplete execution of a transaction T_0 that accesses t-objects X and Y, followed by the step contention-free executions of two transactions T_1 and T_2 which access X and Y respectively. Transactions T_1 and T_2 may contend on a base object since there is a path between X and Y in $G(T_1, T_2, E)$. However, a strict DAP TM implementation would preclude transactions T_1 and T_2 from contending on the same base object since $Dset(T_1) \cap Dset(T_2) = \emptyset$ in E. □

Weak disjoint-access parallelism. Informally, weak DAP means that two transactions can *concurrently contend* on a common base object only if their data sets are connected in the *conflict graph*, capturing data-set overlaps among all concurrent transactions.

Let $G(T_i, T_j, E)$ be an undirected graph whose vertex set is $\bigcup_{T \in \tau_E(T_i, T_j)} Dset(T)$ and there is an edge between t-objects X and Y *iff* there exists $T \in \tau_E(T_i, T_j)$ such that $\{X, Y\} \in Dset(T)$. We say that T_i and T_j are *disjoint-access* in E if there is no path between a t-object in $Dset(T_i)$ and a t-object in $Dset(T_j)$ in $G(T_i, T_j, E)$. A TM implementation M is *weak disjoint-access parallel (weak DAP)* if, for all executions E of M, transactions T_i and T_j concurrently contend on the same base object in E only if T_i and T_j are not disjoint-access in E or there exists a t-object $X \in Dset(T_i) \cap Dset(T_j)$ [24, 110].

We now prove an auxiliary lemma, inspired by [24], concerning weak DAP TM implementations that will be useful in subsequent proofs. Intuitively, the lemma states that, two transactions that are disjoint-access and running one after the other in an execution of a weak DAP TM cannot contend on the same base object.

Lemma 2.10. *Let M be any weak DAP TM implementation. Let $\alpha \cdot \rho_1 \cdot \rho_2$ be any execution of M where ρ_1 (resp., ρ_2) is the step contention-free execution fragment of transaction $T_1 \notin txns(\alpha)$ (resp., $T_2 \notin txns(\alpha)$) and transactions T_1, T_2 are disjoint-access in $\alpha \cdot \rho_1 \cdot \rho_2$. Then, T_1 and T_2 do not contend on any base object in $\alpha \cdot \rho_1 \cdot \rho_2$.*

Proof. Suppose, by contradiction that T_1 and T_2 contend on the same base object in $\alpha \cdot \rho_1 \cdot \rho_2$.

If in ρ_1, T_1 performs a nontrivial event on a base object on which they contend, let e_1 be the last event in ρ_1 in which T_1 performs such an event to some base object b and e_2, the first event in ρ_2 that accesses b. Otherwise, T_1 only performs trivial events in ρ_1 to base objects on which it contends with T_2 in $\alpha \cdot \rho_1 \cdot \rho_2$: let e_2 be the first event in ρ_2 in which ρ_2 performs a nontrivial event to some base object b on which they contend and e_1, the last event of ρ_1 in T_1 that accesses b.

Let ρ_1' (and resp. ρ_2') be the longest prefix of ρ_1 (and resp. ρ_2) that does not include e_1 (and resp. e_2). Since before accessing b, the execution is step contention-free for T_1, $\alpha \cdot \rho_1' \cdot \rho_2'$ is an execution of M. By construction, T_1 and T_2 are disjoint-access in $\alpha \cdot \rho_1' \cdot \rho_2'$ and $\alpha \cdot \rho_1 \cdot \rho_2'$ is indistinguishable to T_2 from $\alpha \cdot \rho_1' \cdot \rho_2'$. Hence, T_1 and T_2 are poised to apply contending events e_1 and e_2 on b in the configuration after $\alpha \cdot \rho_1' \cdot \rho_2'$—a contradiction since T_1 and T_2 cannot concurrently contend on the same base object. □

We now show that weak DAP is a weaker property than RW DAP.

Proposition 2.3. *Weak DAP $\ll$ RW DAP.*

Proof. Clearly, every implementation that satisfies RW DAP also satisfies weak DAP since the conflict graph $\tilde{G}(T_i, T_j, E)$ (for RW DAP) is a subgraph of $G(T_i, T_j, E)$ (for weak DAP).

However, the converse is not true (Algorithm 5.2 describes a TM implementation that satisfies weak DAP, but not RW DAP). Consider the following execution E of a weak DAP TM implementaton M that begins with the t-incomplete execution of a transaction T_0 that reads X and writes to Y, followed by the step contention-free executions of two transactions T_1 and T_2 which write to X and read Y respectively. Transactions T_1 and T_2 may contend on a base object since there is a path between X and Y in $G(T_1, T_2, E)$. However, a RW DAP TM implementation would preclude transactions T_1 and T_2 from contending on the same base object: there is no edge between t-objects X and Y in the corresponding conflict graph $\tilde{G}(T_1, T_2, E)$ because X and Y are not contained in the write set of T_0. □

Thus, the above propositions imply:

Corollary 2.11. *Weak DAP $\ll$ RW DAP $\ll$ Strict DAP $\ll$ Strict data-partitioning.*

2.7 TM complexity metrics

We now present an overview of some of the TM complexity metrics we consider in the thesis.

Step complexity. The step complexity metric, is the total number of events that a process performs on the shared memory, in the worst case, in order to complete its operation on the implementation.

RAW/AWAR patterns. Attiya *et al.* identified two common expensive synchronization patterns that frequently arise in the design of concurrent algorithms: *read-after-write (RAW) or atomic write-after-read (AWAR)* [17, 103] and showed that it is impossible to derive RAW/AWAR-free implementations of a wide class of data types that include *sets*, *queues* and *deadlock-free mutual exclusion.*

Note the shared memory model in the thesis makes the assumption that CPU *events* are performed atomically: every "read" of a base object returns the value of "latest write" to the base object. In practice however, the CPU architecture's *memory model* [3] that specifies the outcome of CPU instructions is *relaxed* without enforcing a strict order among the shared memory instructions. Intuitively, RAW (read-after-write) or AWAR (atomic-write-after-read) patterns [17] capture the amount of "expensive synchronization", *i.e.*, the number of costly memory barriers or conditional primitives [3] incurred by the implementation in relaxed CPU architectures. The metric appears to be more practically relevant than simply counting the number of steps performed by a process, as it accounts for expensive cache-coherence operations or instructions like compare-and-swap. Detailed coverage on memory fences and the RAW/AWAR metric can be found in [103].

Definition 2.16 (Read-after-write metric)**.** *A* RAW *(read-after-write) pattern performed by a transaction T_k in an execution π is a pair of its events e and e', such that: (1) e is a write to a base object b by T_k, (2) e' is a subsequent read of a base object $b' \neq b$ by T_k, and (3) no event on b by T_k takes place between e and e'.*

In the thesis, we are concerned only with *non-overlapping* RAWs, *i.e.*, the read performed by one RAW precedes the write performed by the other RAW.

Definition 2.17 (Atomic write-after-read metric)**.** *An* AWAR *(atomic-write-after-read) pattern e in an execution $\pi \cdot e$ is a nontrivial rmw event on an object b which atomically returns the value of b (resulting after π) and updates b with a new value.*

For example, consider the execution $\pi \cdot e$ where e is the application of a *compare-and-swap* rmw primitive that returns *true*.

Stall complexity. Intuitively, the stall metric captures the fact that the time a process might have to spend before it applies a primitive on a base object can be proportional to the number of processes that try to update the object concurrently.

Let M be any TM implementation. Let e be an event applied by process p to a base object b as it performs a transaction T during an execution E of M. Let $E = \alpha \cdot e_1 \cdots e_m \cdot e \cdot \beta$ be an execution of M, where α and β are execution fragments and $e_1 \cdots e_m$ is a maximal sequence of $m \geq 1$ consecutive nontrivial events by distinct distinct processes other than p that access b. Then, we say that T incurs *m memory stalls in E on account of e*. The *number of memory stalls incurred by T in E* is the sum of memory stalls incurred by all events of T in E [16, 46].

In the thesis, we adopt the following definition of a *k-stall execution* from [16, 46].

Definition 2.18. *An execution $\alpha \cdot \sigma_1 \cdots \sigma_i$ is a k-stall execution for t-operation op executed by process p if*

- *α is p-free,*
- *there are distinct base objects $b_1, \ldots, b_i$ and disjoint sets of processes $S_1, \ldots, S_i$ whose union does not include p and has cardinality k such that, for $j = 1, \ldots i$,*
  - *each process in S_j has an enabled nontrivial event about to access base object b_j after α, and*
  - *in σ_j, p applies events by itself until it is the first about to apply an event to b_j, then each of the processes in S_j applies an event that accesses b_j, and finally, p applies an event that accesses b_j,*
- *p invokes exactly one t-operation op in the execution fragment $\sigma_1 \cdots \sigma_i$*
- *$\sigma_1 \cdots \sigma_i$ contains no events of processes not in $(\{p\} \cup S_1 \cup \cdots \cup S_i)$*
- *in every $(\{p\} \cup S_1 \cup \cdots \cup S_i)$-free execution fragment that extends α, no process applies a nontrivial event to any base object accessed in $\sigma_1 \cdots \sigma_i$.*

Observe that in a k-stall execution E for t-operation *op*, the number of memory stalls incurred by *op* in E is k.

The following lemma will be of use in our proofs.

Lemma 2.12. *Let $\alpha \cdot \sigma_1 \cdots \sigma_i$ be a k-stall execution for t-operation op executed by process p. Then, $\alpha \cdot \sigma_1 \cdots \sigma_i$ is indistinguishable to p from a step contention-free execution [16].*

Remote memory references(RMR) [22]. Modern shared memory CPU architectures employ a *memory hierarchy* [73]: a hierarchy of memory devices with different capacities and costs. Some of the memory is *local* to a given process while the rest of the memory is *remote*. Accesses to memory locations (*i.e.* base objects) that are *remote* to a given process are often orders of magnitude slower than a *local* access of the base object. Thus, the performance of concurrent implementations in the shared memory model may depend on the number of *remote memory references* made to base objects [14].

In the *cache-coherent (CC) shared memory*, each process maintains *local* copies of shared base objects inside its cache, whose consistency is ensured by a *coherence protocol*. Informally, we say that an access to a base object b is *remote* to a process p and causes a *remote memory reference (RMR)* if p's cache contains a cached copy of the object that is out of date or *invalidated*; otherwise the access is *local*.

In the *write-through (CC) protocol*, to read a base object b, process p must have a cached copy of b that has not been invalidated since its previous read. Otherwise, p incurs a RMR. To write to b, p causes a RMR that invalidates all cached copies of b and writes to the main memory.

In the *write-back (CC) protocol*, p reads a base object b without causing a RMR if it holds a cached copy of b in shared or exclusive mode; otherwise the access of b causes a RMR that (1) invalidates all copies of b held in exclusive mode, and writing b back to the main memory, (2) creates a cached copy of b in shared mode. Process p can write to b without causing a RMR if it holds a copy of b in exclusive mode; otherwise p causes a RMR that invalidates all cached copies of b and creates a cached copy of b in exclusive mode.

In the *distributed shared memory (DSM)*, each base object is forever assigned to a single process and it is *remote* to the others. Any access of a remote register causes a RMR.

3

Safety for transactional memory

Arthur: If I asked you where the hell we were, would I regret it?
Ford: We're safe.
Arthur: Oh good.
Ford: We're in a small galley cabin in one of the spaceships of the Vogon Constructor Fleet.
Arthur: Ah, this is obviously some strange use of the word safe that I wasn't previously aware of.

Douglas Adams-The Hitchhiker's Guide to the Galaxy

3.1 Overview

In the context of Transactional memory, intermediate states witnessed by the read operations of an incomplete transaction may affect the user application through the outcome of its read operations. If the intermediate state is not consistent with any sequential execution, the application may experience a fatal irrecoverable error or enter an infinite loop. Thus, it is important that *each* transaction, including *aborted* ones observes a *consistent* state so that the implementation does not export any pathological executions.

A state should be considered consistent if it could result from a serial application of transactions observed in the current execution. In this sense, every transaction should witness a state that *could have been* observed in *some* execution of the sequential code put by the programmer within the transactions. Additionally, a consistent state should not depend on a transaction that has not started committing yet (referred to as *deferred-update* semantics). This restriction appears desirable, since the ongoing transaction may still abort (explicitly by the user or because of consistency reasons) and, thus, render the read inconsistent. Further, the set of histories specified by the consistency criterion must constitute a *safety property*, as defined by Owicki and Lamport [108], Alpern and Schneider [11] and refined by Lynch [100]: it must be non-empty, *prefix-closed* and *limit-closed*.

In this chapter, we define the notion of deferred-update semantics formally, which we then apply to a spectrum of TM consistency criteria. Additionally, we verify if the resulting TM consistency criterion is a safety property, as defined by Lynch [100].

We begin by considering the popular criterion of *opacity* [64], which was the first TM consistency criterion that was proposed to grasp this semantics formally. Opacity requires the states observed by all transactions, included uncommitted ones, to be consistent with a global *serialization*, *i.e.*, a serial execution constituted by committed transactions. Moreover, the serialization should respect the *real-time order*: a transaction that completed before (in real time) another transaction started should appear first in the serialization.

By definition, opacity reduces correctness of a history to correctness of all its prefixes, and thus is prefix-closed and limit-closed by definition. Thus, to verify that a history is opaque, one needs to verify that each of its prefixes is consistent with some global serialization. To simplify verification and explicitly introduce deferred-update semantics into a TM correctness criterion, we specify a general criterion of *du-opacity* [18], which requires the global serial execution to respect the deferred-update property. Informally, a du-opaque history must be indistinguishable from a totally-ordered history, with respect to which no transaction reads from a transaction that has not started committing.

We show that du-opacity is *prefix-closed*, that is, every prefix of a du-opaque history is also du-opaque. We then show that extending opacity (and du-opacity) to infinite histories in a non-trivial way (*i.e.*, requiring that even infinite histories should have proper serializations), does not result in a limit-closed property. However, under certain restrictions, we show that du-opacity is *limit-closed*. In particular, assuming that in an infinite history, every transaction completes each of the operations it invoked, the limit of any sequence of ever extending du-opaque histories is also du-opaque. Therefore, under this assumption, du-opacity is a *safety property* [11, 100, 108], and to prove that a TM implementation that complies with the assumption is du-opaque, it suffices to prove that all its *finite* histories are du-opaque.

One may notice that the intended safety semantics does not require that all transactions observe the same serial execution. Intuitively, to avoid pathological executions, we only need that every transaction witnesses *some* consistent state, while the views of different aborted and incomplete transactions do not have to be consistent with *the same* serial execution. As long as committed transactions constitute a serial execution and every transaction witnesses a consistent state, the execution can be considered "safe": no run-time error that cannot occur in a serial execution can happen. Several definitions like *virtual-world consistency (VWC)* [85] and *Transactional Memory Specification 1 (TMS1)*[43] have adopted this approach. We introduce "deferred-update" versions of these properties and discuss how the resulting properties relate to du-opacity.

Finally, we also study the consistency criterion *Transactional Memory Specification 2 (TMS2)* [43, 98], which was proposed as a restriction of opacity and verify if it is a safety property.

Roadmap of Chapter 3. In Section 3.2 of this chapter, we formally define safety properties. In Section 3.3, we introduce the notion of deferred-update semantics and apply it to the correctness criterions of opacity and strict serializability in Sections 3.4 and 3.5 respectively. Section 3.6 studies two relaxations of opacity: VWC and TMS1 and a restriction of opacity, TMS2. Section 3.7 summarizes the relations between the TM correctness properties proposed in the thesis and presents our concluding remarks.

3.2 Safety properties

A *property* $\mathcal{P}$ is a set of (transactional) histories. Intuitively, a *safety property* says that "no bad thing ever happens".

Definition 3.1 (Lynch [100])**.** *A property $\mathcal{P}$ is a* safety *property if it satisfies the following two conditions:*

Prefix-closure: *For every history $H \in \mathcal{P}$, every prefix H' of H (*i.e.*, every prefix of the sequence of the events in H) is also in $\mathcal{P}$.*

Limit-closure: *For every infinite sequence of finite histories $H^0, H^1, \ldots$ such that for every i, $H^i \in \mathcal{P}$ and H^i is a prefix of H^{i+1}, the limit of the sequence is also in $\mathcal{P}$.*

Notice that the set of histories produced by a TM implementation M is, by construction, prefix-closed. Therefore, every infinite history of M is the limit of an infinite sequence of ever-extending finite histories of M. Thus, to prove that M satisfies a safety property P, it is enough to show that all finite histories of M are in P. Indeed, limit-closure of P then implies that every infinite history of M is also in P.

3.3 Opacity and deferred-update(DU) semantics

In this section, we formalize the notion of deferred-update semantics and apply to the TM correctness condition of *opacity* [64].

Definition 3.2 (Guerraoui and Kapalka [64])**.** *A finite history H is* final-state opaque *if there is a legal t-complete t-sequential history S, such that*

1. *for any two transactions $T_k, T_m \in txns(H)$, if $T_k \prec_H^{RT} T_m$, then $T_k <_S T_m$, and*
2. *S is equivalent to a completion of H.*

We say that S is a final-state serialization *of H.*

Final-state opacity is not prefix-closed. Figure 3.1 depicts a t-complete sequential history H that is final-state opaque, with $T_1 \cdot T_2$ being a legal t-complete t-sequential history equivalent to H. Let $H' = write_1(X,1), read_2(X)$ be a prefix of H in which T_1 and T_2 are t-incomplete. Transaction T_i $(i = 1, 2)$ is completed by inserting $tryC_i \cdot A_i$ immediately after the last event of T_i in H. Observe that neither $T_1 \cdot T_2$ nor $T_2 \cdot T_1$ allow us to derive a serialization of H' (we assume that the initial value of X is 0).

A restriction of final-state opacity, which we refer to as *opacity* [64] explicitly filters out histories that are not prefix-closed.

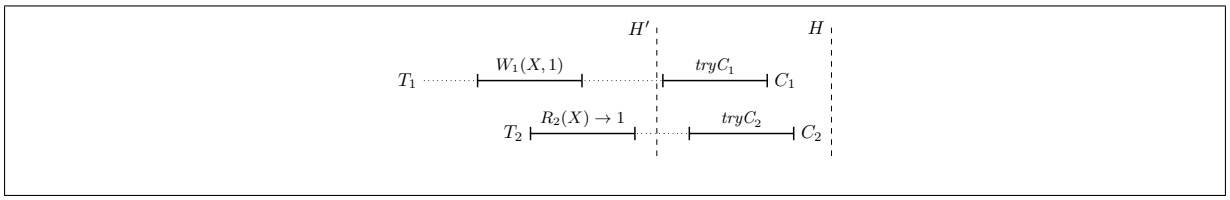

Figure 3.1: History H is final-state opaque, while its prefix H' is not final-state opaque.

Definition 3.3 (Guerraoui and Kapalka [64])**.** *A history H is* opaque *if and only if every finite prefix H' of H (including H itself if it is finite) is final-state opaque.*

It can be easily seen that opacity is prefix- and limit-closed, and, thus, it is a safety property.

We now give a formal definition of opacity with deferred-update semantics. Then we show that the property is prefix-closed and, under certain *liveness* restrictions, limit-closed.

Let H be any history and let S be a legal t-complete t-sequential history that is equivalent to some completion of H. Let $<_S$ be the total order on transactions in S.

Definition 3.4 (Local serialization)**.** *For any $read_k(X)$ that does not return A_k, let $S^{k,X}$ be the prefix of S up to the response of $read_k(X)$ and $H^{k,X}$ be the prefix of H up to the response of $read_k(X)$.*
$S_H^{k,X}$, the local serialization *of $read_k(X)$ with respect to H and S, is the subsequence of $S^{k,X}$ derived by removing from $S^{k,X}$ the events of all transactions $T_m \in txns(H) \setminus \{T_k\}$ such that $H^{k,X}$ does not contain an invocation of $tryC_m()$.*

We are now ready to present our correctness condition, *du-opacity*.

Definition 3.5 (Du-opacity)**.** *A history H is* du-opaque *if there is a legal t-complete t-sequential history S such that*

1. *there is a completion of H that is equivalent to S, and*
2. *for every pair of transactions $T_k, T_m \in txns(H)$, if $T_k \prec_H^{RT} T_m$, then $T_k <_S T_m$, i.e., S respects the real-time ordering of transactions in H, and*
3. *each $read_k(X)$ in S that does not return A_k is legal in $S_H^{k,X}$.*

We then say that S is a (du-opaque) serialization *of H.*

Informally, a history H is du-opaque if there is a legal t-sequential history S that is equivalent to H, respects the real-time ordering of transactions in H and every t-read is legal in its local serialization with respect to H and S. The third condition reflects the implementation's deferred-update semantics, i.e., the legality of a t-read in a serialization does not depend on transactions that start committing after the response of the t-read.

For any du-opaque serialization S, $seq(S)$ denotes the *sequence of transactions* in S and $seq(S)[k]$ denotes the k^{th} transaction in this sequence.

3.4 On the safety of du-opacity

In this section, we examine the safety properties of du-opacity, *i.e.*, whether it is prefix-closed and limit-closed.

3.4.1 Du-opacity is prefix-closed

Lemma 3.1. *Let H be a du-opaque history and let S be a serialization of H. For any $i \in \mathbb{N}$, there is a serialization S^i of H^i (the prefix of H consisting of the first i events), such that $seq(S^i)$ is a subsequence of $seq(S)$.*

Proof. Given H, S and H^i, we construct a t-complete t-sequential history S^i as follows:

- for every transaction T_k that is t-complete in H^i, $S^i|k = S|k$.
- for every transaction T_k that is complete but not t-complete in H^i, $S^i|k$ consists of the sequence of events in $H^i|k$, immediately followed by $tryC_k() \cdot A_k$.
- for every transaction T_k with an incomplete t-operation, $op_k = read_k \vee write_k \vee tryA_k()$ in H^i, $S^i|k$ is the sequence of events in $S|k$ up to the invocation of op_k, immediately followed by A_k.
- for every transaction $T_k \in txns(H^i)$ with an incomplete t-operation, $op_k = tryC_k()$, $S^i|k = S|k$.

By the above construction, S^i is indeed a t-complete history and every transaction that appears in S^i also appears in S. We order transactions in S^i so that $seq(S^i)$ is a subsequence of $seq(S)$.

Note that S^i is derived from events contained in some completion $\overline{H}$ of H that is equivalent to S and some other events to derive a completion of S^i. Since S^i contains events from every complete t-operation in H^i and other events included satisfy Definition 2.1, there is a completion of H^i that is equivalent to S^i.

We now argue that S^i is a serialization of H^i. First we observe that S^i respects the real-time order of H^i. Indeed, if $T_j \prec_{H^i}^{RT} T_k$, then $T_j \prec_H^{RT} T_k$ and $T_j <_S T_k$. Since $seq(S^i)$ is a subsequence of $seq(S)$, we have $T_j <_{S^i} T_k$.

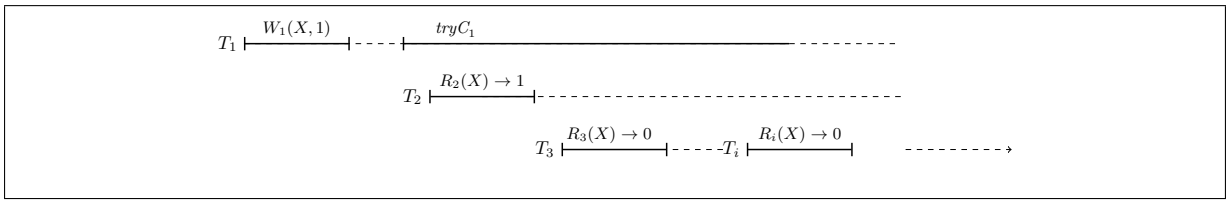

Figure 3.2: An infinite history in which $tryC_1$ is incomplete and any two transactions are concurrent. Each finite prefix of the history is du-opaque, but the infinite limit of the ever-extending sequence is not du-opaque.

To show that S^i is legal, suppose, by way of contradiction, that there is some $read_k(X)$ that returns $v \neq A_k$ in H^i such that v is not the latest written value of X in S^i. If T_k contains a $write_k(X, v')$ preceding $read_k(X)$ such that $v \neq v'$ and v is not the latest written value for $read_k(X)$ in S^i, it is also not the latest written value for $read_k(X)$ in S, which is a contradiction. Thus, the only case to consider is when $read_k(X)$ should return a value written by another transaction.

Since S is a serialization of H, there is a committed transaction T_m that performs the last $write_m(X, v)$ that precedes $read_k(X)$ in T_k in S. Moreover, since $read_k(X)$ is legal in the local serialization of $read_k(X)$ in H with respect to S, the prefix of H up to the response of $read_k(X)$ must contain an invocation of $tryC_m()$. Thus, $read_k(X) \not\prec_H^{RT} tryC_m()$ and $T_m \in txns(H^i)$. By construction of S^i, $T_m \in txns(S^i)$ and T_m is committed in S^i.

We have assumed, towards a contradiction, that v is not the latest written value for $read_k(X)$ in S^i. Hence, there is a committed transaction T_j that performs $write_j(X, v'); v' \neq v$ in S^i such that $T_m <_{S^i} T_j <_{S^i} T_k$. But this is not possible since $seq(S^i)$ is a subsequence of $seq(S)$.

Thus, S^i is a legal t-complete t-sequential history equivalent to some completion of H^i. Now, by the construction of S^i, for every $read_k(X)$ that does not return A_k in S^i, we have ${S^i}_{H^i}^{k,X} = S_H^{k,X}$. Indeed, the transactions that appear before T_k in ${S^i}_{H^i}^{k,X}$ are those with a $tryC$ event before the response of $read_k(X)$ in H and are committed in S. Since $seq(S^i)$ is a subsequence of $seq(S)$, we have ${S^i}_{H^i}^{k,X} = S_H^{k,X}$. Thus, $read_k(X)$ is legal in ${S^i}_{H^i}^{k,X}$. □

Lemma 3.1 implies that every prefix of a du-opaque history has a du-opaque serialization and thus:

Corollary 3.2. *Du-opacity is a prefix-closed property.*

3.4.2 The limit of du-opaque histories

We observe, however, that du-opacity is, in general, not limit-closed. We present an infinite history that is not du-opaque, but each of its prefixes is.

Proposition 3.1. *Du-opacity is not a limit-closed property.*

Proof. Let H^j denote a finite prefix of H of length j. Consider an infinite history H that is the limit of the histories H^j defined as follows (see Figure 3.2):

- Transaction T_1 performs a $write_1(X, 1)$ and then invokes $tryC_1()$ that is incomplete in H.
- Transaction T_2 performs a $read_2(X)$ that overlaps with $tryC_1()$ and returns 1.
- There are infinitely many transactions T_i, $i \geq 3$, each of which performing a single $read_i(X)$ that returns 0 such that each T_i overlaps with both T_1 and T_2.

We now prove that, for all $j \in \mathbb{N}$, H^j is a du-opaque history. Clearly, H^0 and H^1 are du-opaque histories. For all $j > 1$, we first derive a completion of H^j as follows:

1. $tryC_1()$ (if it is contained in H^j) is completed by inserting C_1 immediately after its invocation,
2. for all $i \geq 2$, any incomplete $read_i(X)$ that is contained in H^j is completed by inserting A_i and $tryC_i \cdot A_i$ immediately after its invocation, and
3. for all $i \geq 2$ and every complete $read_j(X)$ that is contained in H^j, we include $tryC_i \cdot A_i$ immediately after the response of this $read_j(X)$.

We can now derive a t-complete t-sequential history S^j equivalent to the above derived completion of H^j from the sequence of transactions $T_3, \ldots, T_i, T_1, T_2$ (depending on which of these transactions participate in H^j), where $i \geq 3$. It is easy to observe that S^j so derived is indeed a serialization of H^j.

However, there is no serialization of H. Suppose that such a serialization S exists. Since every transaction that participates in H must participate in S, there exists $n \in \mathbb{N}$ such that $seq(S)[n] = T_1$. Consider the transaction at index $n+1$, say T_i in $seq(S)$. But for any $i \geq 3$, T_i must precede T_1 in any serialization (by legality), which is a contradiction. □

Notice that all finite prefixes of the infinite history depicted in Figure 3.2 are also opaque. Thus, if we extend the definition of opacity to cover infinite histories in a non-trivial way, i. e., by explicitly defining opaque serializations for infinite histories, we can reformulate Proposition 3.1 for opacity.

3.4.3 Du-opacity is limit-closed for complete histories

We show now that du-opacity is limit-closed if the only infinite histories we consider are those in which every transaction eventually completes (but not necessarily t-completes).

We first prove an auxiliary lemma on du-opaque serializations. For a transaction $T \in txns(H)$, the *live set of T in H*, denoted $Lset_H(T)$ (T included), is defined as follows: every transaction $T' \in txns(H)$ such that neither the last event of T' precedes the first event of T in H nor the last event of T precedes the first event of T' in H is contained in $Lset_H(T)$. We say that transaction $T' \in txns(H)$ *succeeds the live set of T* and we write $T \prec_H^{LS} T'$ if in H, for all $T'' \in Lset_H(T)$, T'' is complete and the last event of T'' precedes the first event of T'.

Lemma 3.3. *Let H be a finite du-opaque history and assume $T_k \in txns(H)$ is a complete transaction in H, such that every transaction in $Lset_H(T_k)$ is complete in H. Then there is a serialization S of H, such that for all $T_k, T_m \in txns(H)$, if $T_k \prec_H^{LS} T_m$, then $T_k <_S T_m$.*

Proof. Since H is du-opaque, there is a serialization $\tilde{S}$ of H.

Let S be a t-complete t-sequential history such that $txns(\tilde{S}) = txns(S)$, and $\forall\, T_i \in txns(\tilde{S}) : S|i = \tilde{S}|i$. We now perform the following procedure iteratively to derive $seq(S)$ from $seq(\tilde{S})$. Initially $seq(S) = seq(\tilde{S})$. For each $T_k \in txns(H)$, let $T_\ell \in txns(H)$ denote the earliest transaction in $\tilde{S}$ such that $T_k \prec_H^{LS} T_\ell$. If $T_\ell <_{\tilde{S}} T_k$ (implying T_k is not t-complete), then move T_k to immediately precede T_ℓ in $seq(S)$.

By construction, S is equivalent to $\tilde{S}$ and for all $T_k, T_m \in txns(H)$; $T_k \prec_H^{LS} T_m$, $T_k <_S T_m$ We claim that S is a serialization of H. Observe that any two transactions that are complete in H, but not t-complete are not related by real-time order in H. By construction of S, for any transaction $T_k \in txns(H)$, the set of transactions that precede T_k in $\tilde{S}$, but succeed T_k in S are not related to T_k by real-time order. Since $\tilde{S}$ respects the real-time order in H, this holds also for S.

We now show that S is legal. Consider any $read_k(X)$ performed by some transaction T_k that returns $v \in V$ in S and let $T_\ell \in txns(H)$ be the earliest transaction in $\tilde{S}$ such that $T_k \prec_H^{LS} T_\ell$. Suppose, by contradiction, that $read_k(X)$ is not legal in S. Thus, there is a committed transaction T_m that performs $write_m(X, v)$ in $\tilde{S}$ such that $T_m = T_\ell$ or $T_\ell <_{\tilde{S}} T_m <_{\tilde{S}} T_k$. Note that, by our assumption, $read_k(X) \prec_H^{RT} tryC_\ell()$. Since $read_k(X)$ must be legal in its local serialization with respect to H and $\tilde{S}$, $read_k(X) \not\prec_H^{RT} tryC_m()$. Thus, $T_m \in Lset_H(T_k)$. Therefore $T_m \neq T_\ell$. Moreover, T_m is complete, and since it commits in $\tilde{S}$, it is also t-complete in H and the last event of T_m precedes the first event of T_ℓ in H, i.e., $T_m \prec_H^{RT} T_\ell$. Hence, T_ℓ cannot precede T_m in $\tilde{S}$—a contradiction.

Observe also that since T_k is complete in H but not t-complete, H does not contain an invocation of $tryC_k()$. Thus, the legality of any other transaction is unaffected by moving T_k to precede T_ℓ in S. Thus, S is a legal t-complete t-sequential history equivalent to some completion of H. The above arguments also prove that every t-read in S is legal in its local serialization with respect to H and S and, thus, S is a serialization of H. □

The proof uses König's Path Lemma [87] formulated as follows. Let G on a rooted directed graph and let v_0 be the root of G. We say that v_k, a vertex of G, is *reachable* from v_0, if there is a sequence of vertices $v_0 \ldots, v_k$ such that for each i, there is an edge from v_i to v_{i+1}. G is *connected* if every vertex in G is reachable from v_0. G is *finitely branching* if every vertex in G has a finite out-degree. G is *infinite* if it has infinitely many vertices.

Lemma 3.4 (König's Path Lemma [87])**.** *If G is an infinite connected finitely branching rooted directed graph, then G contains an infinite sequence of distinct vertices $v_0, v_1, \ldots$, such that v_0 is the* root*, and for every $i \geq 0$, there is an edge from v_i to v_{i+1}.*

Theorem 3.5. *Under the restriction that in any infinite history H, every transaction $T_k \in txns(H)$ is complete, du-opacity is a limit-closed property.*

Proof. We want to show that the limit H of an infinite sequence of finite ever-extending du-opaque histories is du-opaque. By Corollary 3.2, we can assume the sequence of du-opaque histories to be $H^0, H^1, \ldots H^i, H^{i+1}, \ldots$ such that for all $i \in \mathbb{N}$, H^{i+1} is the one-event extension of H^i.

We construct a rooted directed graph G_H as follows:

1. The root vertex of G_H is (H^0, S^0) where S^0 and H^0 contain the initial transaction T_0.
2. Each non-root vertex of G_H is a tuple (H^i, S^i), where S^i is a du-opaque serialization of H^i that satisfies the condition specified in Lemma 3.3: for all $T_k, T_m \in txns(H)$; $T_k \prec^{LS}_{H^i} T_m$ implies $T_k <_{S^i} T_m$. Note that there exist several possible serializations for any H^i. For succinctness, in the rest of this proof, when we refer to a specific S^i, it is understood to be associated with the prefix H^i of H.
3. Let $cseq_i(S^j)$, $j \geq i$, denote the subsequence of $seq(S^j)$ restricted to transactions whose last event in H is a response event and it is contained in H^i. For every pair of vertices $v = (H^i, S^i)$ and $v' = (H^{i+1}, S^{i+1})$ in G_H, there is an edge from v to v' if $cseq_i(S^i) = cseq_i(S^{i+1})$.

The out-degree of a vertex $v = (H^i, S^i)$ in G_H is defined by the number of possible serializations of H^{i+1}, bounded by the number of possible permutations of the set $txns(S^{i+1})$, implying that G_H is *finitely branching*.

By Lemma 3.1, given any serialization S^{i+1} of H^{i+1}, there is a serialization S^i of H^i such that $seq(S^i)$ is a subsequence of $seq(S^{i+1})$. Indeed, the serialization S^i of H^i also respects the restriction specified in Lemma 3.3. Since $seq(S^{i+1})$ contains every complete transaction that takes its last step in H in H^i, $cseq_i(S^i) = cseq_i(S^{i+1})$. Therefore, for every vertex (H^{i+1}, S^{i+1}), there is a vertex (H^i, S^i) such that $cseq_i(S^i) = cseq_i(S^{i+1})$. Thus, we can iteratively construct a path from (H^0, S^0) to every vertex (H^i, S^i) in G_H, implying that G_H is *connected*.

We now apply König's Path Lemma (Lemma 3.4) to G_H. Since G_H is an infinite connected finitely branching rooted directed graph, we can derive an infinite sequence of distinct vertices

$$\mathcal{L} = (H^0, S^0), (H^1, S^1), \ldots, (H^i, S^i), \ldots$$

such that $cseq_i(S^i) = cseq_i(S^{i+1})$.

The rest of the proof explains how to use $\mathcal{L}$ to construct a serialization of H. We begin with the following claim concerning $\mathcal{L}$.

Claim 3.6. *For any $j > i$, $cseq_i(S^i) = cseq_i(S^j)$.*

Proof. Recall that $cseq_i(S^i)$ is a prefix of $cseq_i(S^{i+1})$, and $cseq_{i+1}(S^{i+1})$ is a prefix of $cseq_{i+1}(S^{i+2})$. Also, $cseq_i(S^{i+1})$ is a subsequence of $cseq_{i+1}(S^{i+1})$. Hence, $cseq_i(S^i)$ is a subsequence of $cseq_{i+1}(S^{i+2})$. But, $cseq_{i+1}(S^{i+2})$ is a subsequence of $cseq_{i+2}(S^{i+2})$. Thus, $cseq_i(S^i)$ is a subsequence of $cseq_{i+2}(S^{i+2})$. Inductively, for any $j > i$, $cseq_i(S^i)$ is a subsequence of $cseq_j(S^j)$. But $cseq_i(S^j)$ is the subsequence of $cseq_j(S^j)$ restricted to complete transactions in H whose last step is in H^i. Thus, $cseq_i(S^i)$ is indeed equal to $cseq_i(S^j)$. □

Let $f : \mathbb{N} \to txns(H)$ be defined as follows: $f(1) = T_0$. For every integer $k > 1$, let

$$i_k = \min\{\ell \in \mathbb{N} | \forall j > \ell : cseq_\ell(S^\ell)[k] = cseq_j(S^j)[k]\}$$

Then, $f(k) = cseq_{i_k}(S^{i_k})[k]$.

Claim 3.7. *The function f is* total *and* bijective.

Proof. (Totality and surjectivity)

Since each transaction $T \in txns(H)$ is complete in some prefix H^i of H, for each $k \in \mathbb{N}$, there exists $i \in \mathbb{N}$ such that $cseq_i(S^i)[k] = T$. By Claim 3.6, for any $j > i$, $cseq_i(S^i) = cseq_i(S^j)$. Since a transaction that is complete in H^i w.r.t H is also complete in H^j w.r.t H, it follows that for every $j > i$, $cseq_j(S^j)[k'] = T$, with $k' \geq k$. By construction of G_H and the assumption that each transaction is complete in H, there exists $i \in \mathbb{N}$ such that each $T \in Lset_{H^i}(T)$ is complete in H and its last step is in H^i, and T precedes in S^i every transaction whose first event succeeds the last event of each $T' \in Lset_{H^i}(T)$ in H^i. Indeed, this implies that for each $k \in \mathbb{N}$, there exists $i \in \mathbb{N}$ such that $cseq_i(S^i)[k] = T$; $\forall j > i : cseq_j(S^j)[k] = T$.

This shows that for every $T \in txns(H)$, there are $i, k \in \mathbb{N}$; $cseq_i(S^i)[k] = T$, such that for every $j > i$, $cseq_j(S^j)[k] = T$. Thus, for every $T \in txns(H)$, there is k such that $f(k) = T$.

(Injectivity)

If $f(k)$ and $f(m)$ are transactions at indices k, m of the same $cseq_i(S^i)$, then clearly $f(k) = f(m)$ implies $k = m$. Suppose $f(k)$ is the transaction at index k in some $cseq_i(S^i)$ and $f(m)$ is the transaction at index m in some $cseq_\ell(S^\ell)$. For every $\ell > i$ and $k < m$, if $cseq_i(S^i)[k] = T$, then $cseq_\ell(S^\ell)[m] \neq T$ since $cseq_i(S^i) = cseq_i(S^\ell)$. If $\ell > i$ and $k > m$, it follows from the definition that $f(k) \neq f(m)$. Similar arguments for the case when $\ell < i$ prove that if $f(k) = f(m)$, then $k = m$. □

By Claim 3.7, $\mathcal{F} = f(1), f(2), \ldots, f(i), \ldots$ is an infinite sequence of transactions. Let S be a t-complete t-sequential history such that $seq(S) = \mathcal{F}$ and for each t-complete transaction T_k in H, $S|k = H|k$; and for transaction that is complete, but not t-complete in H, $S|k$ consists of the sequence of events in $H|k$, immediately followed by $tryA_k() \cdot A_k$. Clearly, there is a completion of H that is equivalent to S.

Let $\mathcal{F}^i$ be the prefix of $\mathcal{F}$ of length i, and $\widehat{S}^i$ be the prefix of S such that $seq(\widehat{S}^i) = \mathcal{F}^i$.

Claim 3.8. *Let $\widehat{H}_i^j$ be a subsequence of H^j reduced to transactions $T_k \in txns(\widehat{S}^i)$ such that the last event of T_k in H is a response event and it is contained in H^j. Then, for every i, there is j such that $\widehat{S}^i$ is a serialization of $\widehat{H}_i^j$.*

Proof. Let H^j be the shortest prefix of H (from $\mathcal{L}$) such that for each $T \in txns(\widehat{S}^i)$, if $seq(S^j)[k] = T$, then for every $j' > j$, $seq(S^{j'})[k] = T$. From the construction of $\mathcal{F}$, such j and k exist. Also, we observe that $txns(\widehat{S}^i) \subseteq txns(S^j)$ and $\mathcal{F}^i$ is a subsequence of $seq(S^j)$. Using arguments similar to the proof of Lemma 3.1, it follows that $\widehat{S}^i$ is indeed a serialization of $\widehat{H}_i^j$. □

Since H is complete, there is exactly one completion of H, where each transaction T_k that is not t-complete in H is completed with $tryC_k \cdot A_k$ after its last event. By Claim 7.11, the limit t-sequential t-complete history is equivalent to this completion, is legal, respects the real-time order of H, and ensures that every read is legal in the corresponding local serialization. Thus, S is a serialization of H. □

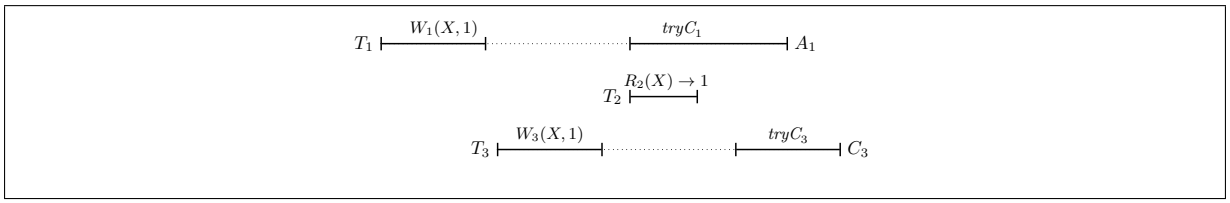

Figure 3.3: A history that is opaque, but not du-opaque.

Theorem 3.5 implies the following:

Corollary 3.9. *Let M be a TM implementation that ensures that in every infinite history H of M, every transaction $T \in txns(H)$ is complete in H. Then, M is du-opaque if and only if every finite history of M is du-opaque.*

3.4.4 Du-opacity vs. opacity

We now compare our deferred-update requirement with the conventional TM correctness property of opacity [64].

Theorem 3.10. *Du-opacity $\subsetneq$ Opacity.*

Proof. We first claim that every finite du-opaque history is opaque. Let H be a finite du-opaque history. By definition, there is a final-state serialization S of H. Since du-opacity is a prefix-closed property, every prefix of H is final-state opaque. Thus, H is opaque.

Again, since every prefix of a du-opaque history is also du-opaque, by Definition 3.3, every infinite du-opaque history is also opaque.

To show that the inclusion is strict, we present an an opaque history that is not du-opaque. Consider the finite history H depicted in Figure 3.3: transaction T_2 performs a $read_2(X)$ that returns the value 1. Observe that $read_2(X) \to 1$ is concurrent to $tryC_1$, but precedes $tryC_3$ in real-time order. Although $tryC_1$ returns A_1 in H, the response of $read_2(X)$ can be justified since T_3 concurrently writes 1 to X and commits. Thus, $read_2(X) \to 1$ *reads-from* transaction T_2 in any serialization of H, but since $read_2(X) \prec_H^{RT} tryC_3$, H is not du-opaque even though each of its prefixes is final-state opaque.

We now formally prove that H is opaque. We proceed by examining every prefix of H.

1. Each prefix up to the invocation of $read_2(X)$ is trivially final-state opaque.
2. Consider the prefix, H^i of H where the i^{th} event is the response of $read_2(X)$. Let S^i be a t-complete t-sequential history derived from the sequence T_1, T_2 by inserting C_1 immediately after the invocation of $tryC_1()$. It is easy to see that S^i is a final-state serialization of H^i.
3. Consider the t-complete t-sequential history S derived from the sequence T_1, T_3, T_2 in which each transaction is t-complete in H. Clearly, S is a final-state serialization of H.

Since H and every (proper) prefix of it are final-state opaque, H is opaque.

Clearly, the required final-state serialization S of H is specified by $seq(S) = T_1, T_3, T_2$ in which T_1 is aborted while T_3 is committed in S (the position of T_1 in the serialization does not affect legality). Consider $read_2(X)$ in S; since $H^{2,X}$, the prefix of H up to the response of $read_2(X)$ does not contain an invocation of $tryC_3()$, the local serialization of $read_2(X)$ with respect to H and S, $S_H^{2,X}$ is $T_1 \cdot read_2(X)$. But $read_2(X)$ is not legal in $S_H^{2,X}$, which is a contradiction. Thus, H is not du-opaque. □

The unique-write case We now show that du-opacity is equivalent to opacity assuming that no two transactions write identical values to the same t-object ("unique-write" assumption).

Let Opacity$_{uw} \subseteq$ Opacity, be a property defined as follows:

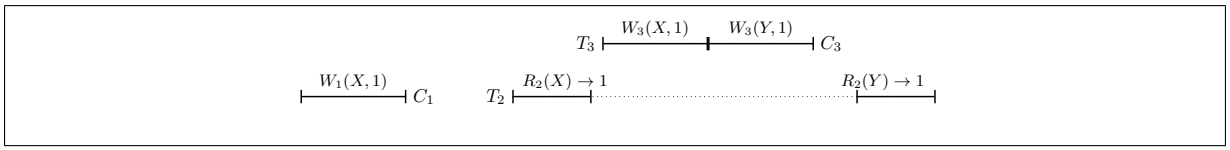

Figure 3.4: A sequential du-opaque history, which is not opaque by the definition of [59].

1. an infinite opaque history $H \in \text{Opacity}_{uw}$ if and only if every transaction $T \in \mathit{txns}(H)$ is complete in H, and

2. an opaque history $H \in \text{Opacity}_{uw}$ if and only if for every pair of write operations $\mathit{write}_k(X, v)$ and $\mathit{write}_m(X, v')$, $v \neq v'$.

Theorem 3.11. *Opacity_{uw} =du-opacity.*

Proof. We show first that every finite history $H \in$Opacity$_{uw}$ is also du-opaque. Let H be any finite opaque history such that for every pair of write operations $\mathit{write}_k(X, v)$ and $\mathit{write}_m(X, v)$, performed by transactions $T_k, T_m \in \mathit{txns}(H)$, respectively, $v \neq v'$.

Since H is opaque, there is a final-state serialization S of H. Suppose by contradiction that H is not du-opaque. Thus, there is a $\mathit{read}_k(X)$ that returns a value $v \in V$ in S that is not legal in $S^{k,X}_H$, the local serialization of $\mathit{read}_k(X)$ with respect to H and S. Let $H^{k,X}$ and $S^{k,X}$ denote the prefixes of H and S, respectively, up to the response of $\mathit{read}_k(X)$ in H and S. Recall that $S^{k,X}_H$, the local serialization of $\mathit{read}_k(X)$ with respect to H and S, is the subsequence of $S^{k,X}$ that does not contain events of any transaction $T_i \in \mathit{txns}(H)$ so that the invocation of $\mathit{tryC}_i()$ is not in $H^{k,X}$. Since $\mathit{read}_k(X)$ is legal in S, there is a committed transaction $T_m \in \mathit{txns}(H)$ that performs $\mathit{write}_m(X, v)$ that is the latest such write in S that precedes T_k. Thus, if $\mathit{read}_k(X)$ is not legal in $S^{k,X}_H$, the only possibility is that $\mathit{read}_k(X) \prec^{RT}_H \mathit{tryC}_m()$. Under the assumption of unique writes, there does not exist any other transaction $T_j \in \mathit{txns}(H)$ that performs $\mathit{write}_j(X, v)$. Consequently, there does not exist any $\overline{H}^{k,X}$ (some completion of $H^{k,X}$) and a t-complete t-sequential history S', such that S' is equivalent to $\overline{H}^{k,X}$ and S' contains any committed transaction that writes v to X. This is, $H^{k,X}$ is not final-state opaque. However, since H is opaque, every prefix of H must be final-state opaque, which is a contradiction.

By Definition 3.3, an infinite history H is opaque if every finite prefix of H is final-state opaque. Theorem 3.5 now implies that $\text{Opacity}_{uw} \subseteq$ du-Opacity.

Definition 3.3 and Corollary 3.2 imply that du-Opacity $\subseteq \text{Opacity}_{uw}$. □

The sequential-history case The deferred-update semantics was mentioned by Guerraoui et al. [59] and later adopted by Kuznetsov and Ravi [90]. In both papers, opacity is only defined for sequential histories, where every invocation of a t-operation is immediately followed by a matching response. In particular, these definitions require the final-state serialization to respect the *read-commit order*: in these definitions, a history H is opaque if there is a final-state serialization S of H such that if a t-read of a t-object X by a transaction T_k precedes the tryC of a transaction T_m that commits on X in H, then T_k precedes T_m in S. As we observed in Figure 3.4, this definition is not equivalent to opacity even for sequential histories.

The property considered in [59, 90] is strictly stronger than du-opacity: the sequential history H in Figure 3.4 is du-opaque (and consequently opaque by Theorem 3.10): a du-opaque serialization (in fact the only possible one) for this history is T_1, T_3, T_2. However, in the restriction of opacity defined above, T_2 must precede T_3 in any serialization, since the response of $\mathit{read}_2(X)$ precedes the invocation of $\mathit{tryC}_3()$.

3.5 Strict serializability with DU semantics

In this section, we discuss the deferred-update restriction of strict serializability from Definition 2.2. First, we remark that, just as final-state opacity, strict serializability is not prefix-closed (cf. Figure 3.1). However, we show that the restriction of deferred-update semantics applied to strict serializability induces a safety property.

Definition 3.6 (Strict serializability with du semantics)**.** *A finite history H is* strictly serializable *if there is a legal t-complete t-sequential history S, such that*

1. *there is a completion $\overline{H}$ of H, such that S is equivalent to $cseq(\overline{H})$, where $cseq(\overline{H})$ is the subsequence of $\overline{H}$ reduced to committed transactions in $\overline{H}$,*
2. *for any two transactions $T_k, T_m \in txns(H)$, if $T_k \prec_H^{RT} T_m$, then T_k precedes T_m in S, and*
3. *each $read_k(X)$ in S that does not return A_k is legal in $S_H^{k,X}$.*

Notice that every du-opaque history is strictly serializable, but not vice-versa.

Theorem 3.12. *Strict serializability is a safety property.*

Proof. Observe that any strictly serializable serialization of a finite history H does not include events of any transaction that has not invoked *tryC* in H.

To show prefix-closure, a proof almost identical to that of Lemma 3.1 implies that, given a strictly serializable history H and a serialization S, there is a serialization S' of H' (H' is some prefix of H) such that $seq(S')$ is a prefix of $seq(S)$.

Consider an infinite sequence of finite histories

$$H^0, \ldots, H^i, H^{i+1}, \ldots,$$

where H^{i+1} is a one-event extension of H^i, we prove that the infinite limit H of this ever-extending sequence is strictly serializable. As in Theorem 3.5, we construct an infinite rooted directed graph G_H: a vertex is a tuple (H^i, S^i) (note that for each $i \in \mathbb{N}$, there are several such vertices of this form), where S^i is a serialization of H^i and there is an edge from (H^i, S^i) to (H^{i+1}, S^{i+1}) if $seq(S^i)$ is a prefix of $seq(S^{i+1})$. The resulting graph is finitely branching since the out-degree of a vertex is bounded by the number of possible serializations of a history. Observe that for every vertex (H^{i+1}, S^{i+1}), there is a vertex $H^i, S^i)$ such that $seq(S^i)$ is a prefix of $seq(S^{i+1})$. Thus, G_H is connected since we can iteratively construct a path from the root (H^0, S^0) to every vertex (H^i, S^i) in G_H. Applying König's Path Lemma to G_H, we obtain an infinite sequence of distinct vertices, $(H^0, S^0), (H^1, S^1), \ldots, (H^i, S^i), \ldots$. Then, $S = \lim_{i \to \infty} S_i$ gives the desired serialization of H. □

3.6 Du-opacity vs. other deferred-update criteria

In this section, we first study two relaxations of opacity: *Virtual-world consistency* [85] and *Transactional Memory Specification 1* [43]. We then study *Transactional Memory Specification 2* which is a restriction of opacity.

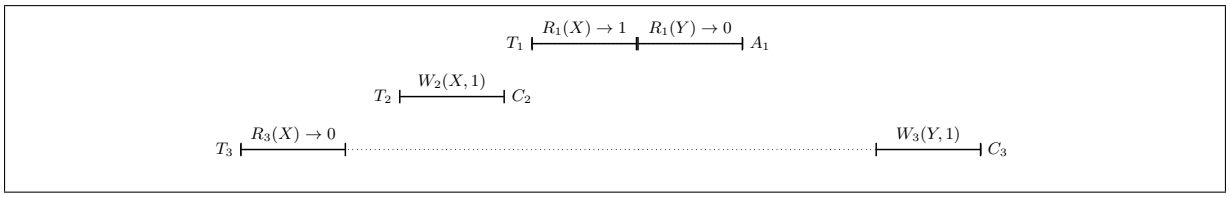

Figure 3.5: A history that is du-VWC, but not du-opaque.

3.6.1 Virtual-world consistency

Virtual World Consistency (VWC) [85] was proposed as a relaxation of opacity (in our case, du-opacity), where each aborted transaction should be consistent with its *causal past* (but not necessarily with a serialization formed by committed transactions). Intuitively, a transaction T_1 causally precedes T_2 if T_2 reads a value written and committed by T_1. The original definition [85] required that no two write operations are ever invoked with the same argument (the *unique-writes* assumption). Therefore, the causal precedence is unambiguously identified for each transactional read. Below we give a more general definition.

Given a t-sequential legal history S and transactions $T_i, T_j \in \mathit{txns}(S)$, we say that T_i *reads* X *from* T_j if (1) T_i reads v in X and (2) T_j is the last committed transaction that writes v to X and precedes T_i in S.

Now consider a (not necessarily t-sequential) history H. We say that T_i *could have read* X *from* T_j in H if T_j writes a value v to a t-object X, T_i reads v in X, and $\mathit{read}_i(X) \not\prec_H^{RT} \mathit{tryC}_j()$.

Given $\mathcal{T} \subseteq \mathit{txns}(H)$, let $H^{\mathcal{T}}$ denote the subsequence of H restricted to events of transactions in $\mathcal{T}$.

Definition 3.7 (du-VWC)**.** *A finite history H is* du-virtual-world consistent *if it is strictly serializable (with du-semantics), and for every aborted or t-incomplete transaction $T_i \in \mathit{txns}(H)$, there is $\mathcal{T} \subseteq \mathit{txns}(H)$ including T_i and a t-sequential t-complete legal history S such that:*

1. *S is equivalent to a completion of $H^{\mathcal{T}}$,*
2. *For all $T_j, T_k \in \mathit{txns}(S)$, if T_j reads X from T_k in S, then T_j could have read X from T_k in H,*
3. *S respects the per-process order of H: if T_j and T_k are executed by the same process and $T_j \prec_H^{RT} T_k$, then $T_j \prec_S T_k$.*

We refer to S as a du-VWC serialization for T_i in H.

Intuitively, with every t-read on X performed by T_i in H, the du-VWC serialization S associates some transaction T_j from which T_i could have read the value of X. Recursively, with every read performed by T_j, S associates some T_m from which T_j could have read, etc. Altogether, we get a "plausible" causal past of T_i that constitutes a serial history. Notice that to ensure deferred-update semantics, we only allow a transaction T_j to read from a transaction T_k that invoked tryC_k by the time of the read operation of T_j.

We now prove that du-VWC is a strictly weaker property than du-opacity. Since du-TMS2 is strictly weaker than du-opacity, it follows that Du-TMS2 $\subsetneq$ du-VWC.

Theorem 3.13. *Du-opacity $\subsetneq$ du-VWC.*

Proof. If a history H is du-opaque, then there is a du-opaque serialization S equivalent to $\overline{H}$, where $\overline{H}$ is some completion of H. By construction, S is a total-order on the set of all transactions that participate in S. Trivially, by taking $\mathcal{T} = \mathit{txns}(H)$, we derive that S is a du-VWC serialization for every aborted or t-incomplete transaction $T_i \in \mathit{txns}(H)$. Indeed, S respects the real-time order and, thus, the per-process order of H. Since S respects the deferred-update order in H, every t-read in S "could have happened" in H.

To show that the inclusion is strict, Figure 3.5 depicts a history H that is du-VWC, but not du-opaque. Clearly, H is strictly serializable. Here T_2, T_1 is the required du-VWC serialization for aborted transaction T_1. However, H has no du-opaque serialization. □ □

Theorem 3.14. *Du-VWC is a safety property.*

Proof. By Definition 3.7, a history H is du-VWC if and only if H is strictly serializable and there is a du-VWC serialization for every transaction $T_i \in \mathit{txns}(H)$ that is aborted or t-incomplete in H.

To prove prefix-closure, recall that strict serializability is a prefix-closed property (Theorem 3.12). Therefore, any du-VWC serialization S for a transaction T_i in history H is also a du-VWC serialization S for a transaction T_i in any prefix of H that contains events of T_i.

To prove limit-closure, consider an infinite sequence of du-VWC histories H^0, H^1, ..., H^i, H^{i+1} , ..., where each H^{i+1} is the one-event extension of H^i and prove that the infinite limit, H of this sequence is also a du-VWC history. Theorem 3.12 establishes that there is a strictly serializable serialization for H.

Since, for all $i \in \mathbb{N}$, H^i is du-VWC, for every transaction T_i that is t-incomplete or aborted in H^i, there is a VWC serialization for T_i. Consequently, there is a du-VWC serialization for every aborted or incomplete transaction T_i in H. □

3.6.2 Transactional memory specification (TMS)

Transactional Memory Specification (TMS) 1 and 2 were formulated in I/O automata [43]. Following [15], we adapt these definitions to our framework and explicitly introduce the deferred-update requirement.

TMS1. Given a history H, TMS1 requires us to justify the behavior of all committed transactions in H by a legal t-complete t-sequential history that preserves the real-time order in H (strict serializability), and to justify the response of each complete t-operation performed in H by a legal t-complete t-sequential history S. The t-sequential history S used to justify a complete t-operation $op_{i,k}$ (the i^{th} t-operation performed by transaction T_k) includes T_k and a subset of transactions from H whose operations justify $op_{i,k}$. (Our description follows [15].)

Let $H^{k,i}$ denote the prefix of a history H up to (and including) the response of i^{th} t-operation $op_{k,i}$ of transaction T_k. We say that a history H'' is a *possible past* of $H^{k,i}$ if H'' is a subsequence of $H^{k,i}$ and consists of all events of transaction T_k and all events from some subset of committed transactions and transactions that have invoked $tryC$ in $H^{k,i}$ such that if a transaction $T \in H''$, then for a transaction $T' \prec^{RT}_{H^{k,i}} T$, $T' \in H''$ if and only if T' is committed in $H^{k,i}$. Let $\mathit{cTMSpast}(H, op_{k,i})$ denote the set of possible pasts of $H^{k,i}$.

For any history $H'' \in \mathit{cTMSpast}(H, op_{k,i})$, let $\mathit{ccomp}(H'')$ denote the history generated from H'' by the following procedure: for all $m \neq k$, replace every event A_m by C_m and complete every incomplete $tryC_m$ with including C_m at the end of H''; include $tryC_k \cdot A_k$ at the end of H''.

Definition 3.8 (du-TMS1)**.** *A history H satisfies* du-TMS1 *if*

1. *H is strictly serializable (with du-semantics), and*
2. *for each complete t-read $op_{i,k}$ that returns a non-A_k response in H, there exist a legal t-complete t-sequential history S and a history H' such that:*
   - *$H' = \mathit{ccomp}(H'')$, where $H'' \in \mathit{cTMSpast}(H, op_{k,i})$*
   - *H' is equivalent to S*
   - *for any two transactions T_k and T_m in H', if $T_k \prec^{RT}_{H'} T_m$ then $T_k <_S T_m$*

We refer to S as the du-TMS1 serialization for $op_{i,k}$.

Theorem 3.15. *Du-TMS1 is a safety property.*

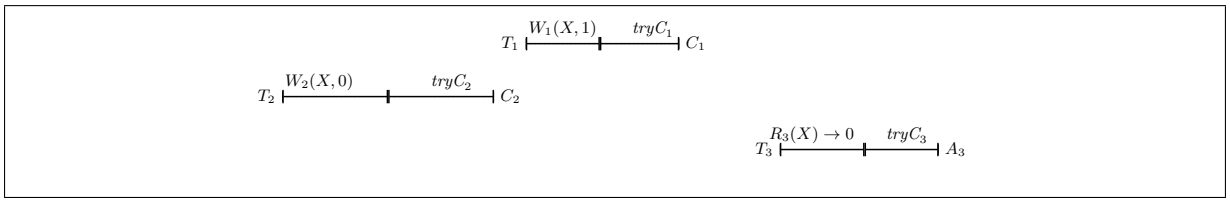

Figure 3.6: A history which is du-VWC but not du-TMS1.

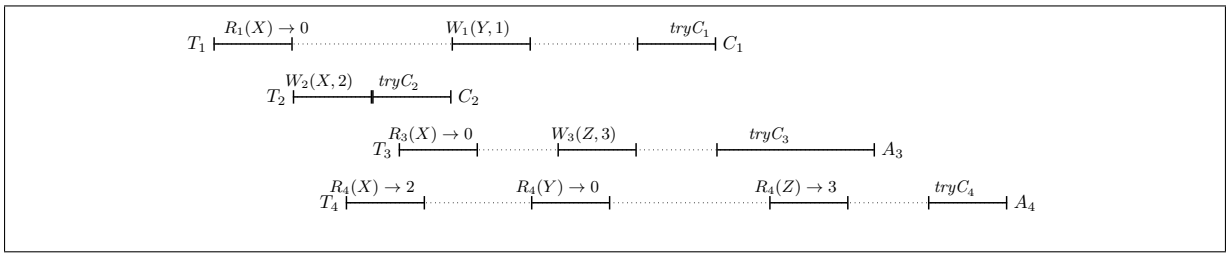

Figure 3.7: A history which is du-TMS1 but not du-VWC.

Proof. A history H is du-TMS1 if and only if H is strictly serializable and there is a du-TMS1 serialization for every t-operation $op_{k,i}$ that does not return A_k in H.

To see that du-TMS1 is prefix closed, recall that strict serializability is a prefix-closed property. Let H be any du-TMS1 history and H^i, any prefix of H. We now need to prove that, for every t-operation $op_{k,i} \neq tryC_k$ that returns a non-A_k response in H^i, there is a du-TMS1 serialization for $op_{k,i}$. But this is immediate since the du-TMS1 serialization for $op_{i,k}$ in H is also the required du-TMS1 serialization for $op_{k,i}$ in H^i.

To see that du-TMS1 is limit closed, consider an infinite sequence

$$H^0, H^1, \ldots H^i, H^{i+1}, \ldots$$

of finite du-TMS1 histories, such that H^{i+1} is a one-event extension of H^i. Let let H be the corresponding infinite limit history. We want to show that H is also du-TMS1.

Since strict serializability is a limit-closed property (Theorem 3.12), H is strictly serializable. By assumption, for all $i \in \mathbb{N}$, H^i is du-TMS1. Thus, for every transaction T_i that participates in H^i, there is a du-TMS1 serialization $S^{i,k}$ for each t-operation $op_{k,i}$. But $S^{i,k}$ is also the required du-TMS1 serialization for $op_{k,1}$ in H. The claim follows. □

It has been shown [98] that Opacity is a strictly stronger property than du-TMS1, that is, Opacity $\subsetneq$ du-TMS1. Since Du-Opacity $\subsetneq$ Opacity (Theorem 3.10) it follows that Du-Opacity $\subsetneq$ du-TMS1. On the other hand, du-TMS1 is incomparable to du-VWC, as demonstrated by the following examples.

Proposition 3.2. *There is a history that is du-TMS1, but not du-VWC.*

Proof. Figure 3.7 depicts a history H that is du-TMS1, but not du-VWC. Observe that H is strictly serializable. To prove that H is du-TMS1, we need to prove that there is a TMS1 serialization for each t-read that returns a non-abort response in H. Clearly, the serialization in which only T_3 participates is the required TMS1 serialization for $read_3(X) \to 0$. Now consider the aborted transaction T_4. The TMS1 serialization for $read_4(X) \to 2$ is T_2, T_4, while the TMS1 serialization that justifies the response of $read_4(Y) - > 0$ includes just T_4 itself. The only nontrivial t-read whose response needs to be justified is $read_4(Z) \to 3$. Indeed, $tryC_3$ overlaps with $read_4(Z)$ and thus, the response of $read_4(Z)$ can be justified by choosing transactions in $cTMSpart(H, read_4(Z))$ to be $\{T_3, T_2, T_4\}$ and then deriving a TMS1 serialization $S = T_3, T_2, T_4$ for $read_4(Z) \to 3$ in which $tryC_3$ may be completed by including the commit response.

However, H is not du-VWC. Consider transaction T_3 which returns A_3 in H: T_3 must be aborted in any serialization equivalent to some direct causal past of T_4. But $read_4(Z)$ returns the value 3 that is written by T_3. Thus, $read_4(Z)$ cannot be legal in any du-VWC serialization for T_4. □

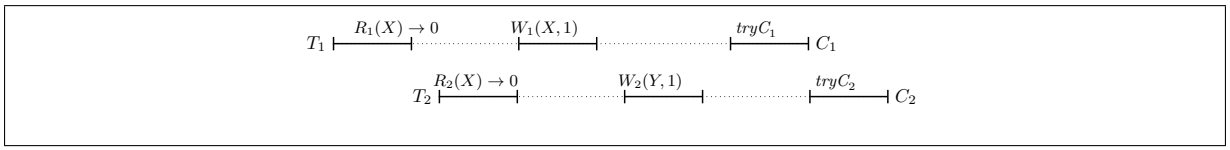

Figure 3.8: A history that is du-opaque, but not TMS2 [43].

Proposition 3.3. *There is a history that is du-VWC, but not du-TMS1.*

Proof. Figure 3.6 depicts a history H that is du-VWC, but not du-TMS1. Clearly, H is strictly serializable. Observe that T_3 could have read only from T_1 in H (T_1 writes the value 0 to X that is returned by $read_3(X)$). Therefore, T_1, T_3 is the required du-VWC serialization for aborted transaction T_3.

However, H is not du-TMS1: since both transactions T_1 and T_2 are committed and precede T_3 in real-time order, they must be included in any du-TMS1 serialization for $read_3(X) \to 0$. But there is no such du-TMS1 serialization that would ensure the legality of $read_3(X)$. □

TMS2. We now study the TMS2 definition which imposes an extra restriction on the opaque serialization.

Definition 3.9 (du-TMS2)**.** *A history H is* du-TMS2 *if there is a legal t-complete t-sequential history S equivalent to some completion, $\overline{H}$ of H such that*

1. *for any two transactions $T_k, T_m \in txns(H)$, such that T_m is a committed updating transaction, if $C_k \prec_H^{RT} tryC_m$ or $A_k \prec_H^{RT} tryC_m$, then $T_k \prec_S T_m$, and*
2. *for any two transactions $T_k, T_m \in txns(H)$, if $T_k \prec_H^{RT} T_m$, then $T_k <_S T_m$, and*
3. *each $read_k(X)$ in S that does not return A_k is legal in $S_H^{k,X}$.*

We refer to S as the du-TMS2 serialization of H.

It has been shown [98] that TMS2 is a strictly stronger property than Opacity, *i.e.*, TMS2 $\subsetneq$ Opacity. We now show that du-TMS2 is strictly stronger than du-opacity. Indeed, from Definition 3.9, we observe that every history that is du-TMS2 is also du-opaque. The following proposition completes the proof.

Proposition 3.4. *There is a history that is du-opaque, but not du-TMS2.*

Proof. Figure 3.8 depicts a history H that is du-opaque, but not du-TMS2. Indeed, there is a du-opaque serialization S of H such that $seq(S) = T_2, T_1$. On the other hand, since T_1 commits before T_2, T_1 must precede T_2 in any du-TMS2 serialization, there does not exist any such serialization that ensures every t-read is legal. Thus, H is not du-TMS2. □

Theorem 3.16. *Du-TMS2 is prefix-closed.*

Proof. Let H be any du-TMS2 history. Then, H is also du-opaque. By Corollary 3.2, for every $i \in \mathbb{N}$, there is a du-opaque serialization S^i for H^i. We now need to prove that, for any two transactions $T_k, T_m \in txns(H^i)$, such that T_m is a committed updating transaction, if $C_k \prec_{H^i}^{RT} tryC_m$ or $A_k \prec_{H^i}^{RT} tryC_m$, there is a du-opaque serialization S^i with the restriction that $T_k \prec_{S^i} T_m$.

Suppose by contradiction that there exist transactions $T_k, T_m \in txns(H^i)$, such that T_m is a committed updating transaction and $C_k \prec_{H^i}^{RT} tryC_m$ or $A_k \prec_{H^i}^{RT} tryC_m$, but T_m must precede T_k in any du-opaque serialization S^i. Since $T_m \not\prec_{H^i}^{RT} T_k$, the only possibility is that T_m performs $write_m(X, v)$ and there is $read_k(X) \to v$. However, by our assumption, $write_k(X, v) \prec_{H^i}^{RT} tryC_m$: thus, $read_k(X)$ is not legal in its local serialization with respect to H^i and S^i—contradicting the assumption that S^i is a du-opaque serialization of H^i. Thus, there is a du-TMS2 serialization for H^i, proving that du-TMS2 is a prefix-closed property. □

	du-opacity	du-VWC	du-TMS1	du-TMS2
du-opacity		$\subsetneq$	$\subsetneq$	$\supsetneq$
du-VWC	$\supsetneq$		$\times$	$\supsetneq$
du-TMS1	$\supsetneq$	$\times$		$\supsetneq$
du-TMS2	$\subsetneq$	$\subsetneq$	$\subsetneq$	

Table 3.1: Relations between TM consistency definitions.

Proposition 3.5. *Du-TMS2 is not limit-closed.*

Proof. The counter-example to establish that du-opacity is not limit-closed (Figure 3.2) also shows that du-TMS2 is not limit-closed: all histories discussed in the counter-example are in du-TMS2. □

3.7 Related work and Discussion

The properties discussed in this chapter explicitly preclude reading from a transaction that has not yet invoked *tryCommit*, which makes them prefix-closed and facilitates their verification. We believe that this constructive definition is useful to TM practitioners, since it streamlines possible implementations of t-read and tryCommit operations.

We showed that du-opacity is limit-closed under the restriction that every operation eventually terminates, while du-VWC and du-TMS1 are (unconditionally) limit-closed, which makes them safety properties [100].

Table 3.1 summarizes the containment relations between the properties discussed in this chapter: opacity, du-opacity, du-VWC, du-TMS1 and du-TMS2. For example, "du-opacity $\subsetneq$ opacity" means that the set of du-opaque histories is a proper subset of the set of opaque histories, *i.e.*, du-opacity is a strictly stronger property than opacity. Incomparable (not related by containment) properties, such as du-TMS1 and du-VWC are marked with $\times$.

Linearizability [27, 83], when applied to objects with *finite nondeterminism* (*i.e.*, an operation applied to a given state may produce only finitely many outcomes) sequential specifications is a safety property [66, 100]. Recently, it has been shown [66] that linearizability is not limit-closed if the implemented object may expose infinite non-determinism [66], that is, an operation applied to a given state may produce infinitely many different outcomes. The limit-closure proof (cf. Theorem 3.5), using König's lemma, cannot be applied with infinite non-determinism, because the out-degree of the graph G_H, constructed for the limit infinite history H, is not finite.

In contrast, the TM abstraction is *deterministic*, since reads and writes behave deterministically in serial executions, yet du-opacity is not limit-closed. It turns out that the graph G_H for the counter-example history H in Figure 3.2 is not connected. For example, one of the finite prefixes of H can be serialized as T_3, T_1, T_2, but no prefix has a serialization T_3, T_1 and, thus, the root is not connected to the corresponding vertex of G_H. Thus, the precondition of König's lemma does not hold for G_H: the graph is in fact an infinite set of isolated vertices. This is because du-opacity requires even incomplete reading transactions, such as T_2, to appear in the serialization, which is not the case for linearizability, where incomplete operations may be removed from the linearization.

4

Complexity bounds for blocking TMs

"I can't believe that!" said Alice
"Can't you?" the Queen said in a pitying tone. "Try again: draw a long breath, and shut your eyes."
Alice laughed. "There's no use trying," she said: "one can't believe impossible things."
"I daresay you haven't had much practice," said the Queen. "When I was your age, I always did it for half-an-hour a day. Why, sometimes I've believed as many as six impossible things before breakfast."

Lewis Carroll-Through the Looking-Glass

4.1 Overview

In this chapter, we present complexity bounds for TM implementations that provide no non-blocking progress guarantees for transactions and typically allow a transaction to *block* (delay) or abort in concurrent executions. We refer to Section 2.7 in Chapter 2 for an overview of the complexity metrics considered in the thesis.

Sequential TMs. We start by presenting complexity bounds for *single-lock* TMs that satisfy sequential TM-progress. We show that a read-only transaction in an opaque TM featured with weak DAP, weak invisible reads, ICF TM-liveness and sequential TM-progress must *incrementally* validate every next read operation. This results in a quadratic (in the size of the transaction's read set) step-complexity lower bound. Secondly, we prove that if the TM-correctness property is weakened to strict serializability, there exist executions in which the tryCommit of some transaction must access a linear (in the size of the transaction's read set) number of distinct base objects. We then show that expensive synchronization

in TMs cannot be eliminated: even single-lock TMs must perform a RAW (read-after-write) or AWAR (atomic-write-after-read) pattern [17].

Progressive TMs. We turn our focus to *progressive* TM implementations which allow a transaction to be aborted only due to read-write conflicts with concurrent transactions. We introduce a new metric called *protected data size* that, intuitively, captures the amount of data that a transaction must exclusively control at some point of its execution. All progressive TM implementations we are aware of (see, *e.g.*, an overview in [62]) use locks or timing assumptions to give an updating transaction exclusive access to all objects in its write set at some point of its execution. For example, lock-based progressive implementations like *TL* [40] and *TL2* [39] require that a transaction grabs all locks on its write set before updating the corresponding base objects. Our result shows that this is an inherent price to pay for providing progressive concurrency: every committed transaction in a progressive and strict DAP TM implementation providing starvation-free TM-liveness must, at some point of its execution, protect every t-object in its write set.

We also present a very cheap progressive opaque strict DAP TM implementation from read-write base objects with constant expensive synchronization and constant memory stall complexities.

Strongly progressive TMs. We then prove that in any *strongly progressive* strictly serializable TM implementation that accesses the shared memory with read, write and conditional primitives, such as compare-and-swap and load-linked/store-conditional, the total number of *remote memory references* (RMRs) that take place in an execution of a progressive TM in which n concurrent processes perform transactions on a single t-object might reach $\Omega(n \log n)$. The result is obtained via a reduction to an analogous lower bound for mutual exclusion [22]. In the reduction, we show that any TM with the above properties can be used to implement a *deadlock-free* mutual exclusion, employing transactional operations on only one t-object and incurring a constant RMR overhead. The lower bound applies to RMRs in both the *cache-coherent (CC)* and *distributed shared memory (DSM)* models, and it appears to be the first RMR complexity lower bound for transactional memory.

We also present a constant expensive synchronization strongly progressive TM implementation from read-write base objects. Our implementation provides *starvation-free* TM-liveness, thus showing one means of circumventing the lower bound of Rachid *et al.* [64] who proved the impossibility of implementing strongly progressive strictly serializable TMs providing *wait-free* TM-liveness from read-write base objects.

Permissive TMs. We conclude our study of blocking TMs by establishing a linear (in the transaction's data set size) separation between the worst-case transaction expensive synchronization complexity of strongly progressive TMs and *permissive* TMs that allow a transaction to abort only if committing it would violate opacity. Specifically, we show that an execution of a transaction in a *permissive opaque* TM implementation that provides starvation-free TM-liveness may require to perform at least one RAW/AWAR pattern *per* t-read.

Roadmap of Chapter 4. Section 4.2 studies "single-lock" TMs that provide minimal progressiveness or sequential TM-progress, Section 4.3 is devoted to progressive TMs while Section 4.4 is on strongly progressive TMs. In Section 4.5, we study the cost of permissive TMs that allow a transaction to abort only if committing it would violate opacity. Finally, we present related work and open questions in Section 4.6.

4.2 Sequential TMs

We begin with "single-lock", *i.e.*, sequential TMs. Our first result proves that a read-only transaction in a sequential TM featured with weak DAP and weak invisible reads must incur the cost of validating its read set. This results in a quadratic (and resp., linear) (in the size of the transaction's read set) step-complexity lower bound if we assume opacity (and resp., strict serializability). Secondly, we show that expensive synchronization cannot be avoided even in such sequential TMs, *i.e.*, a serializable TM must perform a RAW/AWAR even when transactions are guaranteed to commit only in the absence of any concurrency.

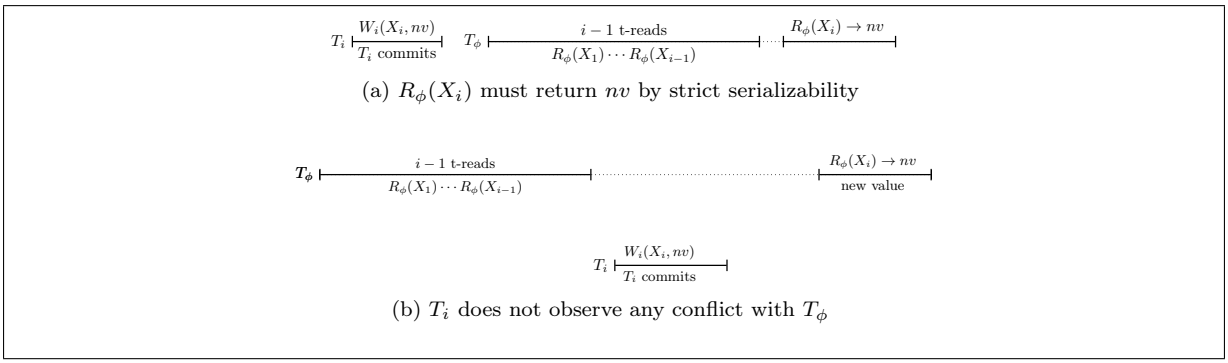

(a) $R_\phi(X_i)$ must return nv by strict serializability

(b) T_i does not observe any conflict with T_ϕ

Figure 4.1: Executions in the proof of Lemma 4.1; By weak DAP, T_ϕ cannot distinguish this from the execution in Figure 4.1a

We first prove the following auxiliary lemma that will be of use in subsequent proofs.

Lemma 4.1. *Let M be any strictly serializable, weak DAP TM implementation that provides sequential TM-progress and sequential TM-liveness. Then, for all $i \in \mathbb{N}$, M has an execution of the form $\pi^{i-1} \cdot \rho^i \cdot \alpha^i$ where,*

- *π^{i-1} is the complete step contention-free execution of read-only transaction T_ϕ that performs $(i-1)$ t-reads: $read_\phi(X_1) \cdots read_\phi(X_{i-1})$,*
- *ρ^i is the t-complete step contention-free execution of a transaction T_i that writes $nv \neq v$ to X_i and commits (v is the initial value of $X - i$),*
- *α_i is the complete step contention-free execution fragment of T_ϕ that performs its i^{th} t-read: $read_\phi(X_i) \to nv_i$.*

Proof. By sequential TM-progress and sequential TM-liveness, M has an execution of the form $\rho^i \cdot \pi^{i-1}$. Since $Dset(T_k) \cap Dset(T_i) = \emptyset$ in $\rho^i \cdot \pi^{i-1}$, by Lemma 2.10, transactions T_ϕ and T_i do not contend on any base object in execution $\rho^i \cdot \pi^{i-1}$. Thus, $\rho^i \cdot \pi^{i-1}$ is also an execution of M.

By assumption of strict serializability, $\rho^i \cdot \pi^{i-1} \cdot \alpha_i$ is an execution of M in which the t-read of X_i performed by T_ϕ must return nv. But $\rho^i \cdot \pi^{i-1} \cdot \alpha_i$ is indistinguishable to T_ϕ from $\pi^{i-1} \cdot \rho^i \cdot \alpha_i$. Thus, M has an execution of the form $\pi^{i-1} \cdot \rho^i \cdot \alpha_i$. □

4.2.1 A quadratic lower bound on step complexity

In this section, we present our step complexity lower bound for sequential TMs.

Theorem 4.2. *For every weak DAP TM implementation M that provides ICF TM-liveness, sequential TM-progress and uses weak invisible reads,*

(1) If M is opaque, for every $m \in \mathbb{N}$, there exists an execution E of M such that some transaction $T \in txns(E)$ performs $\Omega(m^2)$ steps, where $m = |Rset(T_k)|$.

(2) if M is strictly serializable, for every $m \in \mathbb{N}$, there exists an execution E of M such that some transaction $T_k \in txns(E)$ accesses at least $m - 1$ distinct base objects during the executions of the m^{th} t-read operation and $tryC_k()$, where $m = |Rset(T_k)|$.

Proof. For all $i \in \{1, \ldots, m\}$, let v be the initial value of t-object X_i.

(1) Suppose that M is opaque. Let π^m denote the complete step contention-free execution of a transaction T_ϕ that performs m t-reads: $read_\phi(X_1) \cdots read_\phi(X_m)$ such that for all $i \in \{1, \ldots, m\}$, $read_\phi(X_i) \to v$.

By Lemma 4.1, for all $i \in \{2, \ldots, m\}$, M has an execution of the form $E^i = \pi^{i-1} \cdot \rho^i \cdot \alpha_i$.

For each $i \in \{2, \ldots, m\}$, $j \in \{1, 2\}$ and $\ell \leq (i-1)$, we now define an execution of the form $\mathbb{E}^i_{j\ell} = \pi^{i-1} \cdot \beta^\ell \cdot \rho^i \cdot \alpha^i_j$ as follows:

- β^ℓ is the t-complete step contention-free execution fragment of a transaction T_ℓ that writes $nv_\ell \neq v$ to X_ℓ and commits
- α^i_1 (and resp. α^i_2) is the complete step contention-free execution fragment of $read_\phi(X_i) \to v$ (and resp. $read_\phi(X_i) \to A_\phi$).

Claim 4.3. *For all $i \in \{2, \ldots, m\}$ and $\ell \leq (i-1)$, M has an execution of the form $\mathbb{E}^i_{1\ell}$ or $\mathbb{E}^i_{2\ell}$.*

Proof. For all $i \in \{2, \ldots, m\}$, π^{i-1} is an execution of M. By assumption of weak invisible reads and sequential TM-progress, T_ℓ must be committed in $\pi^{i-1} \cdot \rho^\ell$ and M has an execution of the form $\pi^{i-1} \cdot \beta^\ell$. By the same reasoning, since T_i and T_ℓ have disjoint data sets, M has an execution of the form $\pi^{i-1} \cdot \beta^\ell \cdot \rho^i$.

Since the configuration after $\pi^{i-1} \cdot \beta^\ell \cdot \rho^i$ is quiescent, by ICF TM-liveness, $\pi^{i-1} \cdot \beta^\ell \cdot \rho^i$ extended with $read_\phi(X_i)$ must return a matching response. If $read_\phi(X_i) \to v_i$, then clearly $\mathbb{E}^i_1$ is an execution of M with T_ϕ, T_{i-1}, T_i being a valid serialization of transactions. If $read_\phi(X_i) \to A_\phi$, the same serialization justifies an opaque execution.

Suppose by contradiction that there exists an execution of M such that $\pi^{i-1} \cdot \beta^\ell \cdot \rho^i$ is extended with the complete execution of $read_\phi(X_i) \to r$; $r \notin \{A_\phi, v\}$. The only plausible case to analyse is when $r = nv$. Since $read_\phi(X_i)$ returns the value of X_i updated by T_i, the only possible serialization for transactions is T_ℓ, T_i, T_ϕ; but $read_\phi(X_\ell)$ performed by T_k that returns the initial value v is not legal in this serialization—contradiction. □

We now prove that, for all $i \in \{2, \ldots, m\}$, $j \in \{1, 2\}$ and $\ell \leq (i-1)$, transaction T_ϕ must access $(i-1)$ different base objects during the execution of $read_\phi(X_i)$ in the execution $\pi^{i-1} \cdot \beta^\ell \cdot \rho^i \cdot \alpha^i_j$.

By the assumption of weak invisible reads, the execution $\pi^{i-1} \cdot \beta^\ell \cdot \rho^i \cdot \alpha^i_j$ is indistinguishable to transactions T_ℓ and T_i from the execution $\tilde{\pi}^{i-1} \cdot \beta^\ell \cdot \rho^i \cdot \alpha^i_j$, where $Rset(T_\phi) = \emptyset$ in $\tilde{\pi}^{i-1}$. But transactions T_ℓ and T_i are disjoint-access in $\tilde{\pi}^{i-1} \cdot \beta^\ell \cdot \rho^i$ and by Lemma 2.10, they cannot contend on the same base object in this execution.

Consider the $(i-1)$ different executions: $\pi^{i-1} \cdot \beta^1 \cdot \rho^i$, …, $\pi^{i-1} \cdot \beta^{i-1} \cdot \rho^i$. For all $\ell, \ell' \leq (i-1); \ell' \neq \ell$, M has an execution of the form $\pi^{i-1} \cdot \beta^\ell \cdot \rho^i \cdot \beta^{\ell'}$ in which transactions T_ℓ and $T_{\ell'}$ access mutually disjoint data sets. By weak invisible reads and Lemma 2.10, the pairs of transactions $T_{\ell'}$, T_i and $T_{\ell'}$, T_ℓ do not contend on any base object in this execution. This implies that $\pi^{i-1} \cdot \beta^\ell \cdot \beta^{\ell'} \cdot \rho^i$ is an execution of M in which transactions T_ℓ and $T_{\ell'}$ each apply nontrivial primitives to mutually disjoint sets of base objects in the execution fragments β^ℓ and $\beta^{\ell'}$ respectively (by Lemma 2.10).

This implies that for any $j \in \{1, 2\}$, $\ell \leq (i-1)$, the configuration C^i after E^i differs from the configurations after $\mathbb{E}^i_{j\ell}$ only in the states of the base objects that are accessed in the fragment β^ℓ. Consequently, transaction T_ϕ must access at least $i-1$ different base objects in the execution fragment π^i_j to distinguish configuration C^i from the configurations that result after the $(i-1)$ different executions $\pi^{i-1} \cdot \beta^1 \cdot \rho^i$, …, $\pi^{i-1} \cdot \beta^{i-1} \cdot \rho^i$ respectively.

Thus, for all $i \in \{2, \ldots, m\}$, transaction T_ϕ must perform at least $i-1$ steps while executing the i^{th} t-read in π^i_j and T_ϕ itself must perform $\sum_{i=1}^{m-1} i = \frac{m(m-1)}{2}$ steps.

(2) Suppose that M is strictly serializable, but not opaque. Since M is strictly serializable, by Lemma 4.1, it has an execution of the form $E = \pi^{m-1} \cdot \rho^m \cdot \alpha_m$.

For each $\ell \leq (i-1)$, we prove that M has an execution of the form $E_\ell = \pi^{m-1} \cdot \beta^\ell \cdot \rho^m \cdot \bar{\alpha}^m$ where $\bar{\alpha}^m$ is the complete step contention-free execution fragment of $read_\phi(X_m)$ followed by the complete execution of $tryC_\phi$. Indeed, by weak invisible reads, π^{m-1} does not contain any nontrivial events and the execution $\pi^{m-1} \cdot \beta^\ell \cdot \rho^m$ is indistinguishable to transactions T_ℓ and T_m from the executions $\tilde{\pi}^{m-1} \cdot \beta^\ell$ and $\tilde{\pi}^{m-1} \cdot \beta^\ell \cdot \rho^m$ respectively, where $Rset(T_\phi) = \emptyset$ in $\tilde{\pi}^{m-1}$. Thus, applying Lemma 2.10, transactions

$\beta^\ell \cdot \rho^m$ do not contend on any base object in the execution $\pi^{m-1} \cdot \beta^\ell \cdot \rho^m$. By ICF TM-liveness, $read_\phi(X_m)$ and $tryC_\phi$ must return matching responses in the execution fragment $\bar{\alpha}^m$ that extends $\pi^{m-1} \cdot \beta^\ell \cdot \rho^m$. Consequently, for each $\ell \leq (i-1)$, M has an execution of the form $E_\ell = \pi^{m-1} \cdot \beta^\ell \cdot \rho^m \cdot \bar{\alpha}^m$ such that transactions T_ℓ and T_m do not contend on any base object.

Strict serializability of M means that if $read_\phi(X_m) \to nv$ in the execution fragment $\bar{\alpha}^m$, then $tryC_\phi$ must return A_ϕ. Otherwise if $read_\phi(X_m) \to v$ (i.e. the initial value of X_m), then $tryC_\phi$ may return A_ϕ or C_ϕ.

Thus, as with (1), in the worst case, T_ϕ must access at least $m-1$ distinct base objects during the executions of $read_\phi(X_m)$ and $tryC_\phi$ to distinguish the configuration C^i from the configurations after the $m-1$ different executions $\pi^{m-1} \cdot \beta^1 \cdot \rho^m$, ..., $\pi^{m-1} \cdot \beta^{m-1} \cdot \rho^m$ respectively. □

4.2.2 Expensive synchronization in Transactional memory cannot be eliminated

In this section, we show that serializable TMs must perform a RAW/AWAR even if they are guaranteed to commit only when they run in the absence of any concurrency.

Theorem 4.4. *Let M be a serializable TM implementation providing sequential TM-progress and sequential TM-liveness. Then, every execution of M in which a transaction running t-sequentially performs at least one t-read and at least one t-write contains a RAW/AWAR pattern.*

Proof. Consider an execution π of M in which a transaction T_1 running t-sequentially performs (among other events) $read_1(X)$, $write_1(Y, v)$ and $tryC_1()$. Since M satisfies sequential TM-progress and sequential TM-liveness, T_1 must commit in π. Clearly π must contain a write to a base object. Otherwise a subsequent transaction reading Y would return the initial value of Y instead of the value written by T_1.

Let π_w be the first write to a base object in π. Thus, π can be represented as $\pi_s \cdot \pi_w \cdot \pi_f$.

Now suppose by contradiction that π contains neither RAW nor AWAR patterns.

Since π_s contains no writes, the states of base objects in the initial configuration and in the configuration after π_s is performed are the same. Consider an execution $\pi_s \cdot \rho$ where in ρ, a transaction T_2 performs $read_2(Y)$, $write_2(X, 1)$, $tryC_2()$ and commits. Such an execution exists, since ρ is indistinguishable to T_2 from an execution in which T_2 runs t-sequentially and thus T_2 cannot be aborted in $\pi_s \cdot \rho$.

Since π_w contains no AWAR, $\pi_s \cdot \rho \cdot \pi_w$ is an execution of M.

Since $\pi_w \cdot \pi_f$ contains no RAWs, every read performed in $\pi_w \cdot \pi_f$ is applied to base objects which were previously written in $\pi_w \cdot \pi_f$. Thus, there exists an execution $\pi_s \cdot \rho \cdot \pi_w \cdot \pi_f$, such that T_1 cannot distinguish $\pi_s \cdot \pi_w \cdot \pi_f$ and $\pi_s \cdot \rho \cdot \pi_w \cdot \pi_f$. Hence, T_1 commits in $\pi_s \cdot \rho \cdot \pi_w \cdot \pi_f$.

But T_1 reads the initial value of X and T_2 reads the initial value of Y in $\pi_s \cdot \rho \cdot \pi_w \cdot \pi_f$, and thus T_1 and T_2 cannot be both committed (at least one of the committed transactions must read the value written by the other)—a contradiction. □

4.3 Progressive TMs

We move on to the stronger (than sequential TMs) class of progressive TMs. We introduce a new metric called *protected data size* that, intuitively, captures the number of t-objects that a transaction must exclusively control at some prefix of its execution. We first prove that any strict DAP progressive opaque TM must protect its entire write set at some point in its execution.

Secondly, we describe a constant stall, constant RAW/AWAR strict DAP opaque progressive TM that provides invisible reads and is implemented from read-write base objects.

4.3.1 A linear lower bound on the amount of protected data

Let M be a progressive TM implementation providing starvation-free TM-liveness. Intuitively, a t-object X_j is protected at the end of some finite execution π of M if some transaction T_0 is about to atomically change the value of X_j in its next step (*e.g.*, by performing a compare-and-swap) or does not allow any concurrent transaction to read X_j (*e.g.*, by holding a "lock" on X_j).

Formally, let $\alpha \cdot \pi$ be an execution of M such that π is a t-sequential t-complete execution of a transaction T_0, where $Wset(T_0) = \{X_1, \ldots, X_m\}$. Let u_j $(j = 1, \ldots, m)$ denote the value written by T_0 to t-object X_j in π. In this section, let π^t denote the t-th shortest prefix of π. Let π^0 denote the empty prefix.

For any $X_j \in Wset(T_0)$, let T_j denote a transaction that tries to read X_j and commit. Let $E_j^t = \alpha \cdot \pi^t \cdot \rho_j^t$ denote the extension of $\alpha \cdot \pi^t$ in which T_j runs solo until it completes. Note that, since we only require the implementation to be starvation-free, ρ_j^t can be infinite.

We say that $\alpha \cdot \pi^t$ is $(1, j)$-valent if the read operation performed by T_j in $\alpha \cdot \pi^t \cdot \rho_j^t$ returns u_j (the value written by T_0 to X_j). We say that $\alpha \cdot \pi^t$ is $(0, j)$-valent if the read operation performed by T_j in $\alpha \cdot \pi^t \cdot \rho_j^t$ does not abort and returns an "old" value $u \neq u_j$. Otherwise, if the read operation of T_j aborts or never returns in $\alpha \cdot \pi^t \cdot \rho_j^t$, we say that $\alpha \cdot \pi^t$ is $(\bot, j)$-valent.

Definition 4.1. *We say that* T_0 protects *an object* X_j *in* $\alpha \cdot \pi^t$*, where* π^t *is the* t*-th shortest prefix of* π *(*$t > 0$*) if one of the following conditions holds: (1)* $\alpha \cdot \pi^t$ *is* $(0, j)$*-valent and* $\alpha \cdot \pi^{t+1}$ *is* $(1, j)$*-valent, or (2)* $\alpha \cdot \pi^t$ *or* $\alpha \cdot \pi^{t+1}$ *is* $(\bot, j)$*-valent.*

For *strict disjoint-access parallel* progressive TM, we show that every transaction running t-sequentially must protect every t-object in its write set at some point of its execution.

We observe that the no prefix of π can be 0 and 1-valent at the same time.

Lemma 4.5. *There does not exist* π^t*, a prefix of* π*, and* $i, j \in \{1, \ldots, m\}$ *such that* $\alpha \cdot \pi^t$ *is both* $(0, i)$*-valent and* $(1, j)$*-valent.*

Proof. By contradiction, suppose that there exist i, j and $\alpha \cdot \pi^t$ that is both $(0, i)$-valent and $(1, j)$-valent. Since the implementation is strict DAP, there exists an execution of M, $E_{ij}^t = \alpha \cdot \pi^t \cdot \rho_j^t \cdot \rho_i^t$ that is indistinguishable to T_i from $\alpha \cdot \pi^t \cdot \rho_i^t$. In E_{ij}^t, the only possible serialization is T_0, T_j, T_i. But T_i returns the "old" value of X_i and, thus, the serialization is not legal—a contradiction. □

If $\alpha \cdot \pi^t$ is $(0, i)$-valent (resp., $(1, i)$-valent) for some i, we say that it is 0-valent (resp., 1-valent). By Lemma 4.5, the notions of 0-valence and 1-valence are well-defined.

Theorem 4.6. *Let* M *be a progressive, opaque and strict disjoint-access-parallel TM implementation that provides starvation-free TM-liveness. Let* $\alpha \cdot \pi$ *be an execution of* M*, where* π *is a t-sequential t-complete execution of a transaction* T_0*. Then, there exists* π^t*, a prefix of* π*, such that* T_0 *protects* $|Wset(T_0)|$ *t-objects in* $\alpha \cdot \pi^t$*.*

Proof. Let $Wset_{T_0} = \{X_1, \ldots, X_m\}$. Consider two cases:

(1) Suppose that π has a prefix π^t such that $\alpha \cdot \pi^t$ is 0-valent and $\alpha \cdot \pi^{t+1}$ is 1-valent. By Lemma 4.5, there does not exists i, such that $\alpha \cdot \pi^t$ is $(1, i)$-valent and $\alpha \cdot \pi^{t+1}$ is $(0, i)$-valent. Thus, one of the following are true

- For every $i \in \{1, \ldots, m\}$, $\alpha \cdot \pi^t$ is $(0, i)$-valent and $\alpha \cdot \pi^{t+1}$ is $(1, i)$-valent
- At least one of $\alpha \cdot \pi^t$ and $\alpha \cdot \pi^{t+1}$ is $(\bot, i)$-valent, *i.e.*, the t-operation of T_i aborts or never returns

In either case, T_0 protects m t-objects in $\alpha \cdot \pi^t$.

(2) Now suppose that such π^t does not exists, i.e., there is no $i \in \{1, \ldots, m\}$ and $t \in \{0, |\pi| - 1\}$ such that E_i^t exists and returns an old value, and E_i^{t+1} exists and returns a new value.

Suppose there exists s, t, $0 < s + 1 < t$, $S \subseteq \{1, \ldots, m\}$, such that:

- $\alpha \cdot \pi^s$ is 0-valent,
- $\alpha \cdot \pi^t$ is 1-valent,
- for all r, $s < r < t$, and for all $i \in S$, $\alpha \cdot \pi^r$ is $(\perp, i)$-valent.

We say that $s + 1, \ldots, t - 1$ is a *protecting fragment* for t-objects $\{X_j | j \in S\}$.

Since M is opaque and progressive, $\alpha \cdot \pi^0 = \alpha$ is 0-valent and $\alpha \cdot \pi$ is 1-valent. Thus, the assumption of Case (2) implies that for each X_i, there exists a protecting fragment for $\{X_i\}$. In particular, there exists a protecting fragment for $\{X_1\}$.

Now we proceed by induction. Let $\pi_{s+1}, \ldots, \pi_{t-1}$ be a protecting fragment for $\{X_1, \ldots, X_{u-1}\}$ such that $u \leq m$.

Now we claim that there must be a subfragment of $s + 1, \ldots, t - 1$ that protects $\{X_1, \ldots, X_u\}$.

Suppose not. Thus, there exists r, $s < r < t$, such that $\alpha \cdot \pi^r$ is $(0, u)$-valent or $(1, u)$-valent. Suppose first that $\alpha \cdot \pi^r$ is $(1, u)$-valent. Since $\alpha \cdot \pi^s$ is $(0, i)$-valent for some $i \neq u$, by Lemma 4.5 and the assumption of Case (2), there must exist s', t', $s < s' + 1 < t' \leq r$ such that

- $\alpha \cdot \pi^{s'}$ is 0-valent,
- $\alpha \cdot \pi^{t'}$ is 1-valent,
- for all r', $s' < r' < t'$, $\alpha \cdot \pi^{r'}$ is $(\perp, u)$-valent.

As a result, $s' + 1, \ldots, t' - 1$ is a protecting fragment for $\{X_1, \ldots, X_u\}$. The case when $\alpha \cdot \pi^r$ is $(0, u)$-valent is symmetric, except that now we should consider fragment $r, \ldots, t$ instead of $s, \ldots, r$.

Thus, there exists a subfragment of $s + 1, \ldots, t - 1$ that protects $\{X_1, \ldots, X_u\}$. By induction, we obtain a protecting fragment $s'' + 1, \ldots, t'' - 1$ for $\{X_1, \ldots, X_m\}$. Thus, any prefix $\alpha \cdot \pi^r$, where $s'' < r < t''$ protects exactly m t-objects.

In both cases, there is a prefix of $\alpha \cdot \pi$ that protects exactly m t-objects. □

The lower bound of Theorem 4.6 is tight: it is matched by all progressive implementations we are aware of, including Algorithm 4.1 described in the next section.

4.3.2 A constant stall and constant expensive synchronization strict DAP opaque TM

In this section, we describe a cheap progressive, opaque TM implementation LP (Algorithm 4.1). Our TM LP, every transaction performs at most a single RAW, every t-read operation incurs $O(1)$ memory stalls and maintains exactly one version of every t-object at any prefix of an execution. Moreover, the implementation is strict DAP and uses only read-write base objects.

Base objects. For every t-object X_j, LP maintains a base object v_j that stores the *value* of X_j. Additionally, for each X_j, we maintain a bit L_j, which if set, indicates the presence of an updating transaction writing to X_j. Also, for every process p_i and t-object X_j, LP maintains a *single-writer bit* r_{ij} to which only p_i is allowed to write. Each of these base objects may be accessed only via read and write primitives.

Read operations. The implementation first reads the value of t-object X_j from base object v_j and then reads the bit L_j to detect contention with an updating transaction. If L_j is set, the transaction is aborted; if not, read validation is performed on the entire read set. If the validation fails, the transaction

Algorithm 4.1 Strict DAP progressive opaque TM implementation LP; code for T_k executed by process p_i

```
1:  Shared base objects:
2:    v_j, for each t-object X_j, allows reads and writes
3:    r_ij, for each process p_i and t-object X_j
4:        single-writer bit
5:        allows reads and writes
6:    L_j, for each t-object X_j
7:        allows reads and writes
8:  Local variables:
9:    Rset_k, Wset_k for every transaction T_k;
10:       dictionaries storing {X_m, v_m}

11: read_k(X_j):
12:   if X_j ∉ Rset(T_k) then
13:     [ov_j, k_j] := read(v_j)
14:     Rset(T_k) := Rset(T_k) ∪ {X_j, [ov_j, k_j]}
15:     if read(L_j) ≠ 0 then
16:       Return A_k
17:     if validate() then
18:       Return A_k
19:     Return ov_j
20:   else
21:     [ov_j, ⊥] := Rset(T_k).locate(X_j)
22:     Return ov_j

23: write_k(X_j, v):
24:   nv_j := v
25:   Wset(T_k) := Wset(T_k) ∪ {X_j}
26:   Return ok

27: tryC_k():
28:   if |Wset(T_k)| = ∅ then
29:     Return C_k
30:   locked := acquire(Wset(T_k))
31:   if ¬ locked then
32:     Return A_k
33:   if isAbortable() then
34:     release(Wset(T_k))
35:     Return A_k
      // Exclusive write access to each v_j
36:   for all X_j ∈ Wset(T_k) do
37:     write(v_j, [nv_j, k])
38:   release(Wset(T_k))
39:   Return C_k

40: Function: release(Q):
41:   for all X_j ∈ Q do
42:     write(L_j, 0)
43:   for all X_j ∈ Q do
44:     write(r_ij, 0)
45:   Return ok

46: Function: acquire(Q):
47:   for all X_j ∈ Q do
48:     write(r_ij, 1)
49:   if ∃X_j ∈ Q; t ≠ k : read(r_tj) = 1 then
50:     for all X_j ∈ Q do
51:       write(r_ij, 0)
52:     Return false
      // Exclusive write access to each L_j
53:   for all X_j ∈ Q do
54:     write(L_j, 1)
55:   Return true

56: Function: isAbortable() :
57:   if ∃X_j ∈ Rset(T_k) : X_j ∉ Wset(T_k) ∧ read(L_j) ≠ 0
    then
58:     Return true
59:   if validate() then
60:     Return true
61:   Return false

62: Function: validate() :
        // Read validation
63:   if ∃X_j ∈ Rset(T_k):[ov_j, k_j] ≠ read(v_j) then
64:     Return true
65:   Return false
```

is aborted. Otherwise, the implementation returns the value of X_j. For a read-only transaction T_k, $tryC_k$ simply returns the commit response.

Updating transactions. The $write_k(X, v)$ implementation by process p_i simply stores the value v locally, deferring the actual updates to $tryC_k$. During $tryC_k$, process p_i attempts to obtain exclusive write access to every $X_j \in Wset(T_k)$. This is realized through the single-writer bits, which ensure that no other transaction may write to base objects v_j and L_j until T_k relinquishes its exclusive write access to $Wset(T_k)$. Specifically, process p_i writes 1 to each r_{ij}, then checks that no other process p_t has written 1 to any r_{tj} by executing a series of reads (incurring a single RAW). If there exists such a process that concurrently contends on write set of T_k, for each $X_j \in Wset(T_k)$, p_i writes 0 to r_{ij} and aborts T_k. If successful in obtaining exclusive write access to $Wset(T_k)$, p_i sets the bit L_j for each X_j in its write set. Implementation of $tryC_k$ now checks if any t-object in its read set is concurrently contended by another transaction and then validates its read set. If there is contention on the read set or validation fails (indicating the presence of a conflicting transaction), the transaction is aborted. If not, p_i writes the

values of the t-objects to shared memory and relinquishes exclusive write access to each $X_j \in Wset(T_k)$ by writing 0 to each of the base objects L_j and r_{ij}.

Complexity. Read-only transactions do not apply any nontrivial primitives. Any updating transaction performs at most a single RAW in the course of acquiring exclusive write access to the transaction's write set. Thus, every transaction performs $O(1)$ non-overlapping RAWs in any execution.

Recall that a transaction may write to base objects v_j and L_j only after obtaining exclusive write access to t-object X_j, which in turn is realized via single-writer base objects. Thus, no transaction performs a write to any base object b immediately after a write to b by another transaction, *i.e.*, every transaction incurs only $O(1)$ memory stalls on account of any event it performs. The $read_k(X_j)$ implementation reads base objects v_j and L_j, followed by the validation phase in which it reads v_k for each X_k in its current read set. Note that if the first read in the validation phase incurs a stall, then $read_k(X_j)$ aborts. It follows that each t-read incurs $O(1)$ stalls in every execution.

Proof of opacity. We now prove that LP implements an opaque TM.

We introduce the following technical definition: process p_i *holds a lock on* X_j *after an execution* π *of Algorithm 4.1* if π contains the invocation of *acquire(Q)*, $X_j \in Q$ by p_i that returned *true*, but does not contain a subsequent invocation of *release(Q′)*, $X_j \in Q'$, by p_i in π.

Lemma 4.7. *For any object* X_j*, and any execution* π *of Algorithm 4.1, there exists at most one process that* holds *a lock on* X_j *after* π*.*

Proof. Assume, by contradiction, that there exists an execution π after which processes p_i and p_k *hold* a lock on the same object, say X_j. In order to hold the lock on X_j, process p_i writes 1 to register r_{ij} and then checks if any other process p_k has written 1 to r_{kj}. Since the corresponding operation *acquire(Q)*, $X_j \in Q$ invoked by p_i returns *true*, p_i read 0 in r_{kj} in Line 49. But then p_k also writes 1 to r_{kj} and later reads that r_{ij} is 1. This is because p_k can write 1 to r_{kj} only after the read of r_{kj} returned 0 to p_i which is preceded by the write of 1 to r_{ij}. Hence, there exists an object X_j such that $r_{ij} = 1; i \neq k$, but the conditional in Line 49 returns *true* to process p_k— a contradiction. □

Observation 4.8. *Let* π *be any execution of Algorithm 4.1. Then, any updating transaction* $T_k \in txns(\pi)$ *executed by process* p_i *writes to base object* v_j *(in Line 37) for some* $X_j \in Wset(T_k)$ *immediately after* π iff p_i *holds the lock on* X_j *after* π*.*

Lemma 4.9. *Algorithm 4.1 implements an opaque TM.*

Proof. Let E by any finite execution of Algorithm 4.1. Let $<_E$ denote a total-order on events in E.

Let H denote a subsequence of E constructed by selecting *linearization points* of t-operations performed in E. The linearization point of a t-operation *op*, denoted as ℓ_{op} is associated with a base object event or an event performed between the invocation and response of *op* using the following procedure.

Completions. First, we obtain a completion of E by removing some pending invocations and adding responses to the remaining pending invocations involving a transaction T_k as follows: every incomplete $read_k$, $write_k$ operation is removed from E; an incomplete $tryC_k$ is removed from E if T_k has not performed any write to a base object during the *release* function in Line 38, otherwise it is completed by including C_k after E.

Linearization points. Now a linearization H of E is obtained by associating linearization points to t-operations in the obtained completion of E as follows:

- For every t-read op_k that returns a non-A_k value, ℓ_{op_k} is chosen as the event in Line 13 of Algorithm 4.1, else, ℓ_{op_k} is chosen as invocation event of op_k
- For every $op_k = write_k$ that returns, ℓ_{op_k} is chosen as the invocation event of op_k
- For every $op_k = tryC_k$ that returns C_k such that $Wset(T_k) \neq \emptyset$, ℓ_{op_k} is associated with the response of *acquire* in Line 30, else if op_k returns A_k, ℓ_{op_k} is associated with the invocation event of op_k

- For every $op_k = tryC_k$ that returns C_k such that $Wset(T_k) = \emptyset$, ℓ_{op_k} is associated with Line 29

$<_H$ denotes a total-order on t-operations in the complete sequential history H.

Serialization points. The serialization of a transaction T_j, denoted as δ_{T_j} is associated with the linearization point of a t-operation performed within the execution of T_j.

We obtain a t-complete history $\bar{H}$ from H as follows: for every transaction T_k in H that is complete, but not t-complete, we insert $tryC_k \cdot A_k$ after H.

A t-complete t-sequential history S is obtained by associating serialization points to transactions in $\bar{H}$ as follows:

- If T_k is an update transaction that commits, then δ_{T_k} is ℓ_{tryC_k}
- If T_k is a read-only or aborted transaction in $\bar{H}$, δ_{T_k} is assigned to the linearization point of the last t-read that returned a non-A_k value in T_k

$<_S$ denotes a total-order on transactions in the t-sequential history S.

Claim 4.10. *If $T_i \prec_H T_j$, then $T_i <_S T_j$*

Proof. This follows from the fact that for a given transaction, its serialization point is chosen between the first and last event of the transaction implying if $T_i \prec_H T_j$, then $\delta_{T_i} <_E \delta_{T_j}$ implies $T_i <_S T_j$. □

Claim 4.11. *Let T_k be any updating transaction that returns* false *from the invocation of* isAbortable *in Line 33. Then, T_k returns C_k within a finite number of its own steps in any extension of E.*

Proof. Observer that T_k performs the write to base objects v_j for every $X_j \in Wset(T_k)$ and then invokes *release* in Lines 37 and 38 respectively. Since neither of these involve aborting the transaction or contain unbounded loops or waiting statements, it follows that T_k will return C_k within a finite number of its steps. □

Claim 4.12. *S is legal.*

Proof. Observe that for every $read_j(X_m) \to v$, there exists some transaction T_i that performs $write_i(X_m, v)$ and completes the event in Line 37 such that $read_j(X_m) \not\prec_H^{RT} write_i(X_m, v)$. More specifically, $read_j(X_m)$ returns as a non-abort response, the value of the base object v_m and v_m can be updated only by a transaction T_i such that $X_m \in Wset(T_i)$. Since $read_j(X_m)$ returns the response v, the event in Line 13 succeeds the event in Line 37 performed by $tryC_i$. Consequently, by Claim 4.11 and the assignment of linearization points, $\ell_{tryC_i} <_E \ell_{read_j(X_m)}$. Since, for any updating committing transaction T_i, $\delta_{T_i} = \ell_{tryC_i}$, by the assignment of serialization points, it follows that $\delta_{T_i} <_E \delta_{T_j}$.

Thus, to prove that S is legal, it suffices to show that there does not exist a transaction T_k that returns C_k in S and performs $write_k(X_m, v')$; $v' \neq v$ such that $T_i <_S T_k <_S T_j$. Suppose that there exists a committed transaction T_k, $X_m \in Wset(T_k)$ such that $T_i <_S T_k <_S T_j$.

T_i and T_k are both updating transactions that commit. Thus,

$$(T_i <_S T_k) \iff (\delta_{T_i} <_E \delta_{T_k})$$
$$(\delta_{T_i} <_E \delta_{T_k}) \iff (\ell_{tryC_i} <_E \ell_{tryC_k})$$

Since, T_j reads the value of X written by T_i, one of the following is true: $\ell_{tryC_i} <_E \ell_{tryC_k} <_E \ell_{read_j(X_m)}$ or $\ell_{tryC_i} <_E \ell_{read_j(X_m)} <_E \ell_{tryC_k}$. Let T_i and T_k be executed by processes p_i and p_k respectively.

Consider the case that $\ell_{tryC_i} <_E \ell_{tryC_k} <_E \ell_{read_j(X_m)}$.

By the assignment of linearization points, T_k returns a response from the event in Line 30 before the read of v_m by T_j in Line 13. Since T_i and T_k are both committed in E, p_k returns *true* from the event in Line 30 only after T_i writes 0 to r_{im} in Line 44 (Lemma 4.17).

Recall that $read_j(X_m)$ checks if X_m is locked by a concurrent transaction (i.e $L_j \neq 0$), then performs read-validation (Line 15) before returning a matching response. Consider the following possible sequence of events: T_k returns *true* from the *acquire* function invocation, sets L_j to 1 for every $X_j \in Wset(T_k)$ (Line 54) and updates the value of X_m to shared-memory (Line 37). The implementation of $read_j(X_m)$ then reads the base object v_m associated with X_m after which T_k releases X_m by writing 0 to r_{km} and finally T_j performs the check in Line 15. However, $read_j(X_m)$ is forced to return A_j because $X_m \in Rset(T_j)$ (Line 14) and has been invalidated since last reading its value. Otherwise suppose that T_k acquires exclusive access to X_m by writing 1 to r_{km} and returns *true* from the invocation of *acquire*, updates v_m in Line 37), T_j reads v_m, T_j performs the check in Line 15 and finally T_k releases X_m by writing 0 to r_{km}. Again, $read_j(X_m)$ returns A_j since T_j reads that r_{km} is 1—contradiction.

Thus, $\ell_{tryC_i} <_E \ell_{read_j(X)} <_E \ell_{tryC_k}$.

We now need to prove that δ_{T_j} indeed precedes ℓ_{tryC_k} in E.

Consider the two possible cases:

- Suppose that T_j is a read-only or aborted transaction in $\bar{H}$. Then, δ_{T_j} is assigned to the last t-read performed by T_j that returns a non-A$_j$ value. If $read_j(X_m)$ is not the last t-read performed by T_j that returned a non-A$_j$ value, then there exists a $read_j(X_z)$ performed by T_j such that $\ell_{read_j(X_m)} <_E \ell_{tryC_k} <_E \ell_{read_j(X_z)}$. Now assume that ℓ_{tryC_k} must precede $\ell_{read_j(X_z)}$ to obtain a legal S. Since T_k and T_j are concurrent in E, we are restricted to the case that T_k performs a $write_k(X_z, v)$ and $read_j(X_z)$ returns v. However, we claim that this t-read of X_z must abort by performing the checks in Line 15. Observe that T_k writes 1 to L_m, L_z each (Line 54) and then writes new values to base objects v_m, v_z (Line 37). Since $read_j(X_z)$ returns a non-A_j response, T_k writes 0 to L_z before the read of L_z by $read_j(X_z)$ in Line 15. Thus, the t-read of X_z would return A_j (in Line 17 after validation of the read set since X_m has been updated— contradiction to the assumption that it the last t-read by T_j to return a non-A_j response.
- Suppose that T_j is an updating transaction that commits, then $\delta_{T_j} = \ell_{tryC_j}$ which implies that $\ell_{read_j(X_m)} <_E \ell_{tryC_k} <_E \ell_{tryC_j}$. Then, T_j must necessarily perform the checks in Line 33 and read that L_m is 1. Thus, T_j must return A_j—contradiction to the assumption that T_j is a committed transaction.

□

The conjunction of Claims 4.10 and 4.12 establish that Algorithm 4.1 is opaque. □

We can now prove the following theorem:

Theorem 4.13. *Algorithm 4.1 describes a progressive, opaque and strict DAP TM implementation LP that provides wait-free TM-liveness, uses invisible reads, uses only read-write base objects, and for every execution E and transaction $T_k \in txns(E)$:*

- *T_k performs at most a single RAW, and*
- *every t-read operation invoked by T_k incurs $O(1)$ memory stalls in E, and*
- *every complete t-read operation invoked by T_k performs $O(|Rset(T_k)|)$ steps in E.*

Proof. (TM-liveness and TM-progress) Since none of the implementations of the t-operations in Algorithm 4.1 contain unbounded loops or waiting statements, every t-operation op_k returns a matching response after taking a finite number of steps in every execution. Thus, Algorithm 4.1 provides wait-free TM-liveness.

To prove progressiveness, we proceed by enumerating the cases under which a transaction T_k may be aborted.

- Suppose that there exists a $read_k(X_j)$ performed by T_k that returns A_k from Line 15. Thus, there exists a process p_t executing a transaction that has written 1 to r_{tj} in Line 48, but has not yet written 0 to r_{tj} in Line 44 or some t-object in $Rset(T_k)$ has been updated since its t-read by T_k. In both cases, there exists a concurrent transaction performing a t-write to some t-object in $Rset(T_k)$.
- Suppose that $tryC_k$ performed by T_k that returns A_k from Line 31. Thus, there exists a process p_t executing a transaction that has written 1 to r_{tj} in Line 48, but has not yet written 0 to r_{tj} in Line 44. Thus, T_k encounters step-contention with another transaction that concurrently attempts to update a t-object in $Wset(T_k)$.
- Suppose that $tryC_k$ performed by T_k that returns A_k from Line 33. Since T_k returns A_k from Line 33 for the same reason it returns A_k after Line 15, the proof follows.

(Strict disjoint-access parallelism) Consider any execution E of Algorithm 4.1 and let T_i and T_j be any two transactions that participate in E and access the same base object b in E.

- Suppose that T_i and T_j contend on base object v_j or L_j. Since for every t-object X_j, there exists distinct base objects v_j and L_j, T_j and T_j contend on v_j only if $X_j \in Dset(T_i) \cap Dset(T_j)$.
- Suppose that T_i and T_j contend on base object r_{ij}. Without loss of generality, let p_i be the process executing transaction T_i; $X_j \in Wset(T_i)$ that writes 1 to r_{ij} in Line 48. Indeed, no other process executing a transaction that writes to X_j can write to r_{ij}. Transaction T_j reads r_{ij} only if $X_j \in Dset(T_j)$ as evident from the accesses performed in Lines 48, 49, 44, 27.

Thus, T_i and T_j access the same base object only if they access a common t-object.

(Opacity) Follows from Lemma 4.9.

(Invisible reads) Observe that read-only transactions do not perform any nontrivial events. Secondly, in any execution E of Algorithm 4.1, and any transaction $T_k \in txns(E)$, if $X_j \in Rset(T_k)$, T_k does not write to any of the base objects associated with X_j nor write any information that reveals its read set to other transactions.

(Complexity) Consider any execution E of Algorithm 4.1.

- For any $T_k \in txns(E)$, each $read_k$ only applies trivial primitives in E while $tryC_k$ simply returns C_k if $Wset(T_k) = \emptyset$. Thus, Algorithm 4.1 uses invisible reads.
- Any read-only transaction $T_k \in txns(E)$ not perform any RAW or AWAR. An updating transaction T_k executed by process p_i performs a sequence of writes (Line 48 to base objects $\{r_{ij}\} : X_j \in Wset(T_k)$, followed by a sequence of reads to base objects $\{r_{tj}\} : t \in \{1, \ldots, n\}, X_j \in Wset(T_k)$ (Line 49) thus incurring a single multi-RAW.
- Let e be a write event performed by some transaction T_k executed by process p_i in E on base objects v_j and L_j (Lines 37 and 54). Any transaction T_k performs a write to v_j or L_j only after T_k writes 0 to r_{ij}, for every $X_j \in Wset(T_k)$. Thus, by Lemmata 4.17 and 4.9, it follows that events that involve an access to either of these base objects incurs $O(1)$ stalls.

  Let e be a write event on base object r_{ij} (Line 48) while writing to t-object X_j. By Algorithm 4.1, no other process can write to r_{ij}. It follows that any transaction $T_k \in txns(E)$ incurs $O(1)$ memory stalls on account of any event it performs in E. Observe that any t-read $read_k(X_j)$ only accesses base objects v_j, L_j and other value base objects in $Rset(T_k)$. But as already established above, these are $O(1)$ stall events. Hence, every t-read operation incurs $O(1)$-stalls in E.

□

The following corollary follows from Theorems 4.13 and 4.2.

Corollary 4.14. *Let M be any weak DAP progressive opaque TM implementation providing ICF TM-liveness and weak invisible reads. Then, for every execution E and each read-only transaction $T_k \in txns(E)$, T_k performs $\Theta(m^2)$ steps in E, where $m = |Rset_E(T_k)|$.*

Algorithm 4.2 Strict DAP progressive strictly serializable TM implementation; code for T_k executed by process p_i

```
1:  read_k(X_j):
2:    if X_j ∉ Rset(T_k) then
3:      [ov_j, k_j] := read(v_j)
4:      Rset(T_k) := Rset(T_k) ∪ {X_j, [ov_j, k_j]}
5:      if read(L_j) ≠ 0 then
6:        Return A_k
7:      Return ov_j
8:    else
9:      [ov_j, ⊥] := Rset(T_k).locate(X_j)
10:     Return ov_j
```

Similarly, we can prove an almost matching upper bound for Theorem 4.2 for strictly serializable progressive TMs.

Consider Algorithm 4.2 that is a simplification of the opaque progressive TM in Algorithm 4.1: we remove the validation performed in the implementation of a t-read, *i.e.*, Line 17 in Algorithm 4.1; otherwise, the two algorithms are identical. It is easy to see this results in a strictly serializable (but not opaque) TM implementation. Thus,

Theorem 4.15. *Algorithm 4.2 describes a progressive, strictly serializable and strict DAP TM implementation that provides wait-free TM-liveness, uses invisible reads, uses only read-write base objects, and for every execution E and transaction $T_k \in txns(E)$: every t-read operation invoked by T_k performs $O(1)$ steps and $tryC_k$ performs $O(|Rset(T_k)|)$ steps in E.*

Corollary 4.16. *Let M be any weak DAP progressive strictly serializable TM implementation providing ICF TM-liveness and weak invisible reads. Then, for every execution E and each read-only transaction $T_k \in txns(E)$, each $read_k$ performs $O(1)$ steps and $tryC_k$ performs $\Theta(m)$ steps in E, where $m = |Rset_E(T_k)|$.*

4.4 Strongly progressive TMs

In this section, we prove that every strongly progressive strictly serializable TM that uses only read, write and *conditional* primitives has an execution in which in which n concurrent processes perform transactions on a single data item and incur $\Omega(\log n)$ *remote memory references* [14].

We then describe a constant RAW/AWAR strongly progressive TM providing starvation-free TM-liveness from read-write base objects.

4.4.1 A $\Omega(n \log n)$ lower bound on remote memory references

Our lower bound on RMR complexity of strongly progressive TMs is derived by reduction to *mutual exclusion*.

Mutual exclusion. The *mutex object* supports two operations: *Entry* and *Exit*, both of which return the response *ok*. We say that a process p_i *is in the critical section after an execution* π if π contains the invocation of Entry by p_i that returns *ok*, but does not contain a subsequent invocation of Exit by p_i in π.

A mutual exclusion implementation satisfies the following properties:

- (*Mutual-exclusion*) After any execution π, there exists at most one process that is in the critical section.

- (*Deadlock-freedom*) Let π be any execution that contains the invocation of Enter by process p_i. Then, in every extension of π in which every process takes infinitely many steps, some process is in the critical section.
- (*Finite-exit)* Every process completes the Exit operation within a finite number of steps.

We describe an implementation of a mutex object $L(M)$ from a strictly serializable, strongly progressive TM implementation M providing wait-free TM-liveness (Algorithm 4.3). The algorithm is based on the mutex implementation in [84].

Given a sequential implementation, we use a TM to execute the sequential code in a concurrent environment by encapsulating each sequential operation within an *atomic* transaction that replaces each read and write of a t-object with the transactional read and write implementations, respectively. If the transaction commits, then the result of the operation is returned; otherwise if one of the transactional operations aborts. For instance, in Algorithm 4.3, we wish to atomically read a t-object X, write a new value to it and return the old value of X prior to this write. To achieve this, we employ a strictly serializable TM implementation M. Since we assume that M is strongly progressive, in every execution, at least one transaction successfully commits and the value of X is returned.

Shared objects. We associate each process p_i with two alternating identities $[p_i, face_i]$; $face_i \in \{0, 1\}$. The strongly progressive TM implementation M is used to enqueue processes that attempt to enter the critical section within a single t-object X (initially $\perp$). For each $[p_i, face_i]$, $L(M)$ uses a register bit $Done[p_i, face_i]$ that indicates if this face of the process has left the critical section or is executing the Entry operation. Additionally, we use a register $Succ[p_i, face_i]$ that stores the process expected to succeed p_i in the critical section. If $Succ[p_i, face_i] = p_j$, we say that p_j is the *successor of* p_i (and p_i is the *predecessor* of p_j). Intuitively, this means that p_j is expected to enter the critical section immediately after p_i. Finally, $L(M)$ uses a 2-dimensional bit array *Lock*: for each process p_i, there are $n-1$ registers associated with the other processes. For all $j \in \{0, \ldots, n-1\} \setminus \{i\}$, the registers $Lock[p_i][p_j]$ are local to p_i and registers $Lock[p_j][p_i]$ are remote to p_i. Process p_i can only access registers in the *Lock* array that are local or remote to it.

Entry operation. A process p_i adopts a new identity $face_i$ and writes *false* to $Done(p_i, face_i)$ to indicate that p_i has started the Entry operation. Process p_i now initializes the successor of $[p_i, face_i]$ by writing $\perp$ to $Succ[p_i, face_i]$. Now, p_i uses a strongly progressive TM implementation M to atomically store its *pid* and identity i.e., $face_i$ to t-object X and returns the *pid* and identity of its *predecessor*, say $[p_j, face_j]$. Intuitively, this suggests that $[p_i, face_i]$ is scheduled to enter the critical section immediately after $]p_j, face_j]$ exits the critical section. Note that if p_i reads the initial value of t-object X, then it immediately enters the critical section. Otherwise it writes *locked* to the register $Lock[p_i, p_j]$ and sets itself to be the successor of $[p_j, face_j]$ by writing p_i to $Succ[p_j, face_j]$. Process p_i now checks if p_j has started the Exit operation by checking if $Done[p_j, face_j]$ is set. If it is, p_i enters the critical section; otherwise p_i spins on the register $Lock[p_i][p_j]$ until it is *unlocked*.

Exit operation. Process p_i first indicates that it has exited the critical section by setting $Done[p_i, face_i]$, following which it *unlocks* the register $Lock[Succ[p_i, face_i]][p_i]$ to allow p_i's successor to enter the critical section.

Lemma 4.17. *The implementation $L(M)$ (Algorithm 4.3) satisfies mutual exclusion.*

Proof. Let E be any execution of $L(M)$. We say that $[p_i, face_i]$ is the *successor* of $[p_j, face_j]$ if p_i reads the value of *prev* in Line 25 to be $[p_j, face_j]$ (and $[p_j, face_j]$ is the *predecessor* of $[p_i, face_i]$); otherwise if p_i reads the value to be $\perp$, we say that p_i has no predecessor.

Suppose by contradiction that there exist processes p_i and p_j that are both inside the critical section after E. Since p_i is inside the critical section, either (1) p_i read $prev = \perp$ in Line 23, or (2) p_i read that $Done[prev]$ is *true* (Line 29) or p_i reads that $Done[prev]$ is *false* and $Lock[p_i][prev.pid]$ is *unlocked* (Line 30).

(Case 1) Suppose that p_i read $prev = \perp$ and entered the critical section. Since in this case, p_i does not have any predecessor, some other process that returns successfully from the *while* loop in Line 25 must

Algorithm 4.3 Mutual-exclusion object L from a strongly progressive, strict serializable TM M; code for process p_i; $1 \leq i \leq n$

```
1:  Local variables:
2:    bit face_i, for each process p_i

3:  Shared objects:
4:    strongly progressive, strictly
5:      serializable TM M
6:      t-object X, initially ⊥
7:      storing value v ∈ {[p_i, face_i]} ∪ {⊥}
8:    for each tuple [p_i, face_i]
9:      Done[p_i, face_i] ∈ {true, false}
10:     Succ[p_i, face_i] ∈ {p_1, ..., p_n} ∪ {⊥}
11:   for each p_i and j ∈ {1, ..., n} \ {i}
12:     Lock[p_i][p_j] ∈ {locked, unlocked}

13: Function: func():
14:   atomic using M
15:       value := tx-read(X)
16:       tx-write(X, [p_i, face_i])
17:   on abort Return false
18:   Return value

19: Entry:
20:   face_i := 1 − face_i
21:   Done[p_i, face_i].write(false)
22:   Succ[p_i, face_i].write(⊥)
23:   while (prev ← func) = false do
24:     no op
25:   end while
26:   if prev ≠ ⊥ then
27:     Lock[p_i][prev.pid].write(locked)
28:     Succ[prev].write(p_i)
29:     if Done[prev] = false then
30:       while Lock[p_i][prev.pid] = unlocked do
31:         no op
32:       end while
33:   Return ok
34:   // Critical section

35: Exit:
36:   Done[p_i, face_i].write(true)
37:   Lock[Succ[p_i, face_i]][p_i].write(unlocked)
38:   Return ok
```

be successor of p_i in E. Since there exists $[p_j, face_j]$ also inside the critical section after E, p_j reads that either $[p_i, face_i]$ or some other process to be its predecessor. Observe that there must exist some such process $[p_k, face_k]$ whose predecessor is $[p_i, face_i]$. Hence, without loss of generality, we can assume that $[p_j, face_j]$ is the successor of $[p_i, face_i]$. By our assumption, $[p_j, face_j]$ is also inside the critical section. Thus, p_j *locked* the register $Lock[p_j, p_i]$ in Line 27 and set itself to be p_i's successor in Line 28. Then, p_j read that $Done[p_i, face_i]$ is *true* or read that $Done[p_i, face_i]$ is *false* and waited until $Lock[p_j, p_i]$ is *unlocked* and then entered the critical section. But this is possible only if p_i has left the critical section and updated the registers $Done[p_i, face_i]$ and $Lock[p_j, p_i]$ in Lines 36 and 37 respectively—contradiction to the assumption that $[p_i, face_i]$ is also inside the critical section after E.

(Case 2) Suppose that p_i did not read $prev = \perp$ and entered the critical section. Thus, p_i read that $Done[prev]$ is *false* in Line 29 and $Lock[p_i][prev.pid]$ is *unlocked* in Line 30, where *prev* is the predecessor of $[p_i, face_i]$. As with case 1, without loss of generality, we can assume that $[p_j, face_j]$ is the successor of $[p_i, face_i]$ or $[p_j, face_j]$ is the predecessor of $[p_i, face_i]$.

Suppose that $[p_j, face_j]$ is the predecessor of $[p_i, face_i]$, *i.e.*, p_i writes the value $[p_i, face_i]$ to the register $Succ[p_j, face_j]$ in Line 28. Since $[p_j, face_j]$ is also inside the critical section after E, process p_i must read that $Done[p_j, face_j]$ is *true* in Line 29 and $Lock[p_i, p_j]$ is *locked* in Line 30. But then p_i could not have entered the critical section after E—contradiction.

Suppose that $[p_j, face_j]$ is the successor of $[p_i, face_i]$, *i.e.*, p_j writes the value $[p_j, face_j]$ to the register $Succ[p_i, face_i]$. Since both p_i and p_j are inside the critical section after E, process p_j must read that $Done[p_i, face_i]$ is *true* in Line 29 and $Lock[p_j, p_i]$ is *locked* in Line 30. Thus, p_j must spin on the register $Lock[p_j, p_i]$, waiting for it to be *unlocked* by p_i before entering the critical section—contradiction to the assumption that both p_i and p_j are inside the critical section.

Thus, $L(M)$ satisfies mutual-exclusion. □

Lemma 4.18. *The implementation $L(M)$ (Algorithm 4.3) provides deadlock-freedom.*

Proof. Let E be any execution of $L(M)$. Observe that a process may be stuck indefinitely only in Lines 23 and 30 as it performs the *while* loop.

Since M is strongly progressive, in every execution E that contains an invocation of Enter by process p_i, some process returns *true* from the invocation of *func*() in Line 23.

Now consider a process p_i that returns successfully from the *while* loop in Line 23. Suppose that p_i is stuck indefinitely as it performs the *while* loop in Line 30. Thus, no process has *unlocked* the register $Lock[p_i][prev.pid]$ by writing to it in the Exit section. Recall that since $[p_i, face_i]$ has reached the *while* loop in Line 30, $[p_i, face_i]$ necessarily has a predecessor, say $[p_j, face_j]$, and has set itself to be p_j's successor by writing p_i to register $Succ[p_j, face_j]$ in Line 28. Consider the possible two cases: the predecessor of $[p_j, face_j$ is some process $p_k; k \neq i$ or the predecessor of $[p_j, face_j$ is the process p_i itself.

(Case 1) Since by assumption, process p_j takes infinitely many steps in E, the only reason that p_j is stuck without entering the critical section is that $[p_k, face_k]$ is also stuck in the *while* loop in Line 30. Note that it is possible for us to iteratively extend this execution in which p_k's predecessor is a process that is not p_i or p_j that is also stuck in the *while* loop in Line 30. But then the last such process must eventually read the corresponding *Lock* to be *unlocked* and enter the critical section. Thus, in every extension of E in which every process takes infinitely many steps, some process will enter the critical section.

(Case 2) Suppose that the predecessor of $[p_j, face_j$ is the process p_i itself. Thus, as $[p_i, face_]$ is stuck in the *while* loop waiting for $Lock[p_i, p_j]$ to be *unlocked* by process p_j, p_j leaves the critical section, *unlocks* $Lock[p_i, p_j]$ in Line 37 and prior to the read of $Lock[p_i, p_j]$, p_j re-starts the Entry operation, writes *false* to $Done[p_j, 1 - face_j]$ and sets itself to be the successor of $[p_i, face_i]$ and spins on the register $Lock[p_j, p_i]$. However, observe that process p_i, which takes infinitely many steps by our assumption must eventually read that $Lock[p_i, p_j]$ is *unlocked* and enter the critical section, thus establishing deadlock-freedom. □

We say that a TM implementation M *accesses a single t-object* if in every execution E of M and every transaction $T \in txns(E)$, $|Dset(T)| \leq 1$. We can now prove the following theorem:

Theorem 4.19. *Any strictly serializable, strongly progressive TM implementation M that accesses a single t-object implies a* deadlock-free, finite exit *mutual exclusion implementation $L(M)$ such that the RMR complexity of M is within a constant factor of the RMR complexity of $L(M)$.*

Proof. (Mutual-exclusion) Follows from Lemma 4.17.

(Finite-exit) The proof is immediate since the Exit operation contains no unbounded loops or waiting statements.

(Deadlock-freedom) Follows from Lemma 4.18.

(RMR complexity) First, let us consider the CC model. Observe that every event not on M performed by a process p_i as it performs the Entry or Exit operations incurs $O(1)$ RMR cost clearly, possibly barring the *while* loop executed in Line 30. During the execution of this *while* loop, process p_i spins on the register $Lock[p_i][p_j]$, where p_j is the predecessor of p_i. Observe that p_i's cached copy of $Lock[p_i][p_j]$ may be invalidated only by process p_j as it *unlocks* the register in Line 37. Since no other process may write to this register and p_i terminates the *while* loop immediately after the write to $Lock[p_i][p_j]$ by p_j, p_i incurs $O(1)$ RMR's. Thus, the overall RMR cost incurred by M is within a constant factor of the RMR cost of $L(M)$.

Now we consider the DSM model. As with the reasoning for the CC model, every event not on M performed by a process p_i as it performs the Entry or Exit operations incurs $O(1)$ RMR cost clearly, possibly barring the *while* loop executed in Line 30. During the execution of this *while* loop, process p_i spins on the register $Lock[p_i][p_j]$, where p_j is the predecessor of p_i. Recall that $Lock[p_i][p_j]$ is a register that is local to p_i and thus, p_i does not incur any RMR cost on account of executing this loop. It follows that p_i incurs $O(1)$ RMR cost in the DSM model. Thus, the overall RMR cost of M is within a constant factor of the RMR cost of $L(M)$ in the DSM model. □

Theorem 4.20. *([22]) Any deadlock-free, finite-exit mutual exclusion implementation from read, write and conditional primitives has an execution whose RMR complexity is $\Omega(n \log n)$.*

Theorems 4.20 and 4.19 imply:

Theorem 4.21. *Any strictly serializable, strongly progressive TM implementation with wait-free TM-liveness from read, write and conditional primitives that accesses a single t-object has an execution whose RMR complexity is* $\Omega(n \log n)$.

4.4.2 A constant expensive synchronization opaque TM

In this section, we describe a strongly progressive opaque TM implementation providing starvation-free TM-liveness from read-write base objects with constant RAW/AWAR cost. For our implementation, we define and implement a *starvation-free multi-trylock* object.

Starvation-free multi-trylock. A *multi-trylock* provides exclusive write-access to a set Q of t-objects. Specifically, a *multi-trylock* exports the following operations

- *acquire(Q)* returns *true* or *false*
- *release(Q)* releases the lock and returns *ok*
- *isContended*(X_j), $X_j \in Q$ returns *true* or *false*

We assume that processes are well-formed: they never invoke a new operation on the multi-trylock before receiving response from the previous invocation.

We say that a process p_i *holds a lock on* X_j *after an execution* π if π contains the invocation of *acquire(Q)*, $X_j \in Q$ by p_i that returned *true*, but does not contain a subsequent invocation of *release(Q′)*, $X_j \in Q'$, by p_i in π. We say that X_j is *locked after* π by process p_i if p_i holds a lock on X_j after π.

We say that X_j is *contended by* p_i *after an execution* π if π contains the invocation of *acquire(Q)*, $X_j \in Q$, by p_i but does not contain a subsequent return *false* or return of *release(Q′)*, $X_j \in Q'$, by p_i in π.

Let an execution π contain the invocation i_{op} of an operation *op* followed by a corresponding response r_{op} (we say that π *contains op*). We say that X_j *is uncontended (resp., locked) during the execution of op in* π if X_j is uncontended (resp., locked) after every prefix of π that contains i_{op} but does not contain r_{op}.

We implement a *multi-trylock* object whose operations are *starvation-free*. The algorithm is inspired by the *Black-White Bakery Algorithm* [121] and uses a finite number of bounded registers.

A starvation-free multi-trylock implementation satisfies the following properties:

- *Mutual-exclusion*: For any object X_j, and any execution π, there exists at most one process that *holds* a lock on X_j after π.
- *Progress*: Let π be any execution that contains *acquire*(Q) by process p_i. If no other process $p_k, k \neq i$ contends infinitely long on some $X_j \in Q$, then *acquire*(Q) returns *true* in π.
- Let π be any execution that contains *isContended*(X_j) invoked by p_i.
  - If X_j is locked by $p_\ell; \ell \neq i$ during the complete execution of *isContended*(X_j) in π, then *isContended*(X_j) returns *true*.
  - If $\forall \ell \neq i$, X_j is never contended by p_ℓ during the execution of *isContended*(X_j) in π, then *isContended*(X_j) returns *false*.

Our starvation-free multi-trylock in Algorithm 4.4 uses the following shared variables: registers r_{ij} for each process p_i and object X_j, a shared bit $color \in \{B, W\}$, registers $LA_i \in \{0, \ldots, N\}$ for each p_i that denote a *Label* and $MC_i \in \{B, W\}$ for each p_i.

We say $(LA_i, i) < (LA_k, k)$ *iff* $LA_i < LA_k$ or $LA_i = LA_k$ and $i < k$. We now prove the following invariant about the multi-trylock implementation.

Lemma 4.22. *In every execution* π *of Algorithm 4.4, if* p_i holds *a lock on some object* X_j *after* π*, then one of the following conditions must hold:*

Algorithm 4.4 Starvation-free multi-trylock invoked by process p_i

```
1:  Shared variables:
2:    LA_i, for each process p_i, initially 0
3:    MC_i ∈ {B, W} for each process p_i, initially W
4:    color ∈ {B, W}, initally W
5:    r_ij, for each process p_i and each t-object X_j, initially 0

6:  acquire(Q):
7:    for all X_j ∈ Q do
8:      write(r_ij, 1)
9:    c_i := color
10:   write(MC_i, c_i)
11:   write(LA_i, 1 + max({LA_k}|MC_k = MC_i))
12:   while ∃j : ∃k ≠ i: isContended(X_j) && ((LA_k ≠ 0; (MC_k = MC_i); (LA_k, k) < (LA_i, i)) ||
13:    (LA_k ≠ 0; (MC_k ≠ MC_i); MC_i = color)) do
14:     no op
15:   end while
16:   return true

17: release(Q):
18:   for all X_j ∈ Q do
19:     write(r_ij, 0)
20:   if MC_i = B then
21:     write(color, W)
22:   else
23:     write(color, B)
24:   write(LA_i, 0)
25:   return ok

26: isContended(X_j):
27:   if ∃p_t : r_tj ≠ 0, t ≠ i then
28:     return true
29:   return false
```

(1) for some $k \neq i$; $LA_k \neq 0$, if $MC_k = MC_i$, then $(LA_k, k) > (LA_i, i)$

(2) for some $k \neq i$; $LA_k \neq 0$, if $MC_k \neq MC_i$, then $MC_i \neq color$

Proof. In order to hold the lock on X_j, some process p_i writes 1 to r_{ij}, writes a value, say W to MC_i and reads the Labels of other processes that have obtained the same color as itself and generates a Label greater by one than the maximum Label read (Line 11). Observe that until the value of the *color* bit is changed, all processes read the same value W. The first process p_i to hold the lock on X_j changes the *color* bit to B when releasing the lock and hence the value read by all subsequent processes will be B until it is changed again. Now consider two cases:

(1) Assume that there exists a process p_k, $k \neq i$, $LA_k \neq 0$ and $MC_k = MC_i$ such that $(LA_k, k) < (LA_i, i)$, but p_i holds a lock on X_j after π. Thus, $isContended(X_j)$ returns *true* to p_i because p_k writes to r_{kj} (Line 8) before writing to LA_k (Line 11). By assumption, $(LA_k, k) < (LA_i, i); LA_k > 0$ and $MC_i = MC_k$, but the conditional in Line 13 returned *true* to p_i without waiting for p_k to stop contending on X_j—contradiction.

(2) Assume that there exists a process p_k, $k \neq i$, $LA_k \neq 0$ and $MC_k \neq MC_i$ such that $MC_i = color$, but p_i holds a lock on X_j after π. Again, since $LA_k > 0$, $isContended(X_j)$ returns *true* to p_i, $MC_k \neq MC_i$ and $MC_i = color$, but the conditional in Line 13 returned *true* to p_i without waiting for p_k to stop contending on X_j—contradiction.

□

We can thus prove the following theorem:

Theorem 4.23. *Algorithm 4.4 is an implementation of multi-trylock object in which every operation is starvation-free and incurs at most four RAWs.*

Proof. Denote by L the shared object implemented by Algorithm 4.4.

Assume, by contradiction, that L does not provide mutual-exclusion: there exists an execution π after which processes p_i and p_k, $k \neq i$ *hold* a lock on the same object, say X_j. Since both p_i and p_k have performed the write to LA_i and LA_k resp. in Line 11, $LA_i, LA_k > 0$. Consider two cases:

(1) If $MC_k = MC_i$, then from Condition 1 of Lemma 4.22, we have $(LA_k, k) < (LA_i, i)$ and $(LA_k, k) > (LA_i, i)$—contradiction.

(2) If $MC_k \neq MC_i$, then from Condition 2 of Lemma 4.22, we have $MC_i \neq color$ and $MC_k \neq color$ which implies $MC_k = MC_i$—contradiction.

L also ensures progress. If process p_i wants to hold the lock on an object X_j i.e. invokes $acquire(Q), X_j \in Q$, it checks if any other process p_k holds the lock on X_j. If such a process p_k exists and $MC_k = MC_i$, then clearly $isContended(X_j)$ returns *true* for p_i and $(LA_k, k) < (LA_i, i)$. Thus, p_i fails the conditional in Line 13 and waits until p_k releases the lock on X_j to return *true*. However, if p_k contends infinitely long on X_j, p_i is also forced to wait indefinitely to be returned *true* from the invocation of $acquire(Q)$. The same argument works when $MC_k \neq MC_i$ since when p_k stops contending on X_j, $isContended(X_j)$ eventually returns *false* for p_i if p_k does not contend infinitely long on X_j.

All operations performed by L are starvation-free. Each process p_i that successfully holds the lock on an object X_j in an execution π invokes $acquire(Q), X_j \in Q$, obtains a color and chooses a value for LA_i since there is no way to be blocked while writing to LA_i. The response of operation $acquire(Q)$ by p_i is only delayed if there exists a concurrent invocation of $acquire(Q'), X_j \in Q'$ by p_k in π. In that case, process p_i waits until p_k invokes $release(Q)$ and writes 0 to r_{kj} and eventually holds the lock on X_j. The implementation of *release* and *isContended* are wait-free operations (and hence starvation-free) since they contains no unbounded loops or waiting statements.

The implementation of $isContended(X_j)$ only reads base objects. The implementation of $release(Q)$ writes to a series of base objects (Line 18) and then reads a base object (Line 20) incurring a single RAW. The implementation of $acquire(Q)$ writes to base objects (Line 8), reads the shared bit *color* (Line 9)—one RAW, writes to a base object (Line 10), reads the Labels (Line 11)—one RAW, writes to its own Label and finally performs a sequence of reads when evaluating the conditional in Line 13—one RAW.

Thus, Algorithm 4.4 incurs at most four RAWs. □

Strongly progressive TM from starvation-free multi-trylock. We now use the starvation-free multi-trylock to implement a starvation-free strongly progressive opaque TM implementation with constant expensive synchronization (Algorithm 4.5). The implementation is almost identical to the progressive TM implementation LP in Algorithm 4.1, except that the function calls to *acquire* and *release* the transaction's write set are replaced with analogous calls to a multi-trylock object.

Theorem 4.24. *Algorithm 4.5 implements a strongly progressive opaque TM implementation with starvation-free t-operations that uses invisible reads and employs at most four RAWs per transaction.*

Proof. (Opacity) Since Algorithm 4.5 is similar to the opaque progressive TM implementation in Algorithm 4.1, it is easy to adapt the proof of Lemma 4.9 to prove opacity for this implementation.

(TM-progress and TM-liveness) Every transaction T_k in a TM M whose t-operations are defined by Algorithm 4.5 can be aborted in the following scenarios:

- Read-validation failed in $read_k$ or $tryC_k$
- $read_k$ or $tryC_k$ returned A_k because $X_j \in Rset(T_k)$ is locked (belongs to write set of a concurrent transaction)

Algorithm 4.5 Strongly progressive, opaque TM: the implementation of T_k executed by p_i

```
1:  Shared variables:
2:    v_j, for each t-object X_j
3:    L, a starvation-free multi-trylock object
4:  Local variables:
5:    Rset_k, Wset_k for every transaction T_k;
6:        dictionaries storing {X_m, v_m}

7:  read_k(X_j):
8:    if X_j ∉ Rset(T_k) then
9:      [ov_j, k_j] := read(v_j)
10:     Rset(T_k) := Rset(T_k) ∪ {X_j, [ov_j, k_j]}
11:     if isAbortable() then
12:       Return A_k
13:     Return ov_j
14:   else
15:     [ov_j, ⊥] := Rset(T_k).locate(X_j)
16:     Return ov_j

17: write_k(X_j, v):
18:   nv_j := v
19:   Wset(T_k) := Wset(T_k) ∪ {X_j}
20:   Return ok

21: tryC_k():
22:   if |Wset(T_k)| = ∅ then
23:     Return C_k
24:   locked := L.acquire(Wset(T_k))
25:   if isAbortable() then
26:     L.release(Wset(T_k))
27:     Return A_k
28:   for all X_j ∈ Wset(T_k) do
29:     write(v_j, (nv_j, k))
30:   L.release(Wset(T_k))
31:   return C_k

32: Function: isAbortable():
33:   if ∃X_j ∈ Rset(T_k) : X_j ∉ Wset(T_k) ∧
      L.isContended(X_j) then
34:     Return true
35:   if validate() then
36:     Return true
37:   Return false
```

Since in each of these cases, a transaction is aborted only because of a read-write conflict with a concurrent transaction, it is easy to see that M is progressive.

To show Algorithm 4.5 also implements a strongly progressive TM, we need to show that for every set of transactions that concurrently contend on a single t-object, at least one of the transactions is not aborted.

Consider transactions T_i and T_k that concurrently attempt to execute $tryC_i$ and $tryC_k$ such that $X_j \in Wset_i \cup Wset_k$. Consequently, they both invoke the *acquire* operation of the multi-trylock (Line 24) and thus, from Theorem 4.23, both T_i and T_k must commit eventually. Also, if validation of a t-read in T_k fails, it means that the t-object is overwritten by some transaction T_i such that T_i precedes T_k, implying at least one of the transactions commit. Otherwise, if some t-object $X_j \in Rset(T_k)$ is locked and returns *abort* since the t-object is in the write set of a concurrent transaction T_i. While it may still be possible that T_i returns A_i after acquiring the lock on $Wset_i$, strong progressiveness only guarantees progress for transactions that conflict on at most one t-object. Thus, in either case, for every set of transactions that conflict on at most one t-object, at least one transaction is not forcefully aborted.

Starvation-free TM-liveness follows from the fact that the multi try-lock we use in the implementation of M provides starvation-free *acquire* and *release* operations.

(Complexity) Any process executing a transaction T_k holds the lock on $Wset(T_k)$ only once during $tryC_k$. If $|Wset(T_k)| = \emptyset$, then the transaction simply returns C_k incurring no RAW's. Thus, from Theorem 4.23, Algorithm 4.5 incurs at most four RAWs per updating transaction and no RAW's are performed in read-only transactions. □

4.5 On the cost of permissive opaque TMs

We have shown that (strongly) progressive TMs that allow a transaction to be aborted only on read-write conflicts have constant RAW/AWAR complexity. However, not aborting on conflicts may not necessarily affect TM-correctness. Ideally, we would like to derive TM implementations that are *permissive*, in the sense that a transaction is aborted only if committing it would violate TM-correctness.

Definition 4.2 (Permissiveness)**.** *A TM implementation M is* permissive with respect to TM-correctness C *if for every history H of M such that H ends with a response r_k and replacing r_k with some $r_k \neq A_k$ gives a history that satisfies C, we have $r_k \neq A_k$.*

Therefore, permissiveness does not allow a transaction to abort, unless committing it would violate the execution's correctness.

We first show that a transaction in a permissive opaque implementation can only be forcefully aborted if it tries to commit:

Lemma 4.25. *Let a TM implementation M be permissive with respect to opacity. If a transaction T_i is forcefully aborted executing a t-operation op_i, then op_i is $tryC_i$.*

Proof. Suppose, by contradiction, that there exists a history H of M such that some $op_i \in \{read_i, write_i\}$ executed within a transaction T_i returns A_i. Let H_0 be the shortest prefix of H that ends just before op_i returns. By definition, H_0 is opaque and any history $H_0 \cdot r_i$ where $r_i \neq A_i$ is not opaque. Let H'_0 be the serialization of H_0.

If op_i is a write, then $H_0 \cdot ok_i$ is also opaque - no write operation of the incomplete transaction T_i appears in H'_0 and, thus, H'_0 is also a serialization of $H_0 \cdot ok_i$.

If op_i is a $read(X)$ for some t-object X, then we can construct a serialization of $H_0 \cdot v$ where v is the value of X written by the last committed transaction in H'_0 preceding T_i or the initial value of X if there is no such transaction. It is easy to see that H_0" obtained from H'_0 by adding $read(X) \cdot v$ at the end of T_i is a serialization of $H_0 \cdot read(X)$. In both cases, there exists a non-A_i response r_i to op_i that preserves opacity of $H_0 \cdot r_i$, and, thus, the only t-operation that can be forcefully aborted in an execution of M is $tryC$. □

We now show that an execution of a transaction in a *permissive opaque* TM implementation (providing starvation-free TM-liveness) may require to perform at least one RAW/AWAR pattern *per* t-read.

Theorem 4.26. *Let M be a permissive opaque TM implementation providing starvation-free TM-liveness. Then, for any $m \in \mathbb{N}$, M has an execution in which some transaction performs m t-reads such that the execution of each t-read contains at least one RAW or AWAR.*

Proof. Consider an execution E of M consisting of transactions T_1, T_2, T_3 as shown in Figure 4.2: T_3 performs a t-read of X_1, then T_2 performs a t-write on X_1 and commits, and finally T_1 performs a series of reads from objects $X_1, \ldots, X_m$. Since the implementation is permissive, no transaction can be forcefully aborted in E, and the only valid serialization of this execution is T_3, T_2, T_1. Note also that the execution generates a sequential history: each invocation of a t-operation is immediately followed by a matching response. Thus, since we assume starvation-freedom as a liveness property, such an execution exists.

We consider $read_1(X_k)$, $2 \leq k \leq m$ in execution E. Imagine that we modify the execution E as follows. Immediately after $read_1(X_k)$ executed by T_1 we add $write_3(X, v)$, and $tryC_3$ executed by T_3 (let $TC_3(X_k)$ denote the complete execution of $W_3(X_k, v)$ followed by $tryC_3$). Obviously, $TC_3(X_k)$ must return abort: neither T_3 can be serialized before T_1 nor T_1 can be serialized before T_3. On the other hand if $TC_3(X_k)$ takes place just before $read_1(X_k)$, then $TC_3(X_k)$ must return commit but $read_1(X_k)$ must return the value written by T_3. In other words, $read_1(X_k)$ and $TC_3(X_k)$ are *strongly non-commutative* [17]: both of them see the difference when ordered differently. As a result, intuitively, $read_1(X_k)$ needs to perform a RAW or AWAR to make sure that the order of these two "conflicting" operations is properly maintained. We formalize this argument below.

Consider a modification E' of E, in which T_3 performs $write_3(X_k)$ immediately after $read_1(X_k)$ and then tries to commit. In any serialization of E', T_3 must precede T_2 ($read_3(X_1)$ returns the initial value of X_1) and T_2 must precede T_1 to respect the real-time order of transactions. The execution of $read_1(X_k)$ does not modify base objects, hence, T_3 does not observe $read_1(X_k)$ in E'. Since M is permissive, T_3 must commit in E'. But since T_1 performs $read_1(X_k)$ before T_3 commits and T_3 updates X_k, we also have

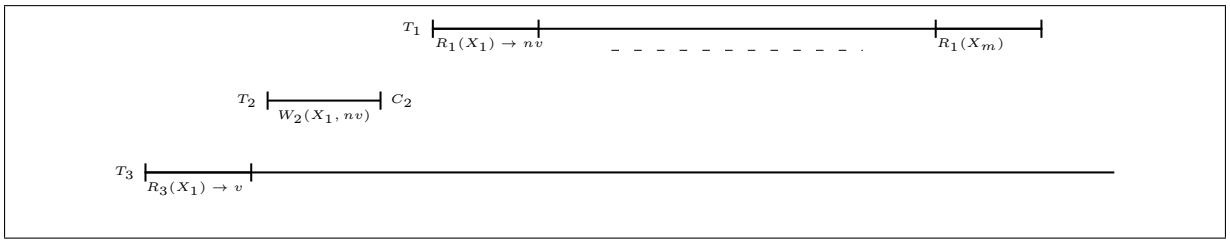

Figure 4.2: Execution E of a permissive, opaque TM: T_2 and T_3 force T_1 to perform a RAW/AWAR in each $R_1(X_k)$, $2 \leq k \leq m$

T_1 must precede T_3 in any serialization. Thus, T_3 cannot precede T_1 in any serialization—contradiction. Consequently, each $read_1(X_k)$ must perform a write to a base object.

Let π be the execution fragment that represents the complete execution of $read_1(X_k)$ and E^k, the prefix of E up to (but excluding) the invocation of $read_1(X_k)$.

Clearly, π contains a write to a base object. Let π_w be the first write to a base object in π. Thus, π can be represented as $\pi_s \cdot \pi_w \cdot \pi_f$. Suppose that π does not contain a RAW or AWAR. Consider the execution fragment $E^k \cdot \pi_s \cdot \rho$, where ρ is the complete execution of $TC_3(X_k)$ by T_3. Such an execution of M exists since π_s does not perform any base object write, hence, $E^k \cdot \pi_s \cdot \rho$ is indistinguishable to T_3 from $E^k \cdot \rho$.

Since, by our assumption, $\pi_w \cdot \pi_f$ contains no RAW, any read performed in $\pi_w \cdot \pi_f$ can only be applied to base objects previously written in $\pi_w \cdot \pi_f$. Since π_w is not an AWAR, $E^k \cdot \pi_s \cdot \rho \cdot \pi_w \cdot \pi_f$ is an execution of M since it is indistinguishable to T_1 from $E^k \cdot \pi$. In $E^k \cdot \pi_s \cdot \rho \cdot \pi_w \cdot \pi_f$, T_3 commits (as in ρ) but T_1 ignores the value written by T_3 to X_k. But there exists no serialization that justifies this execution—contradiction to the assumption that M is opaque. Thus, each $read_1(X_k)$, $2 \leq k \leq m$ must contain a RAW/AWAR.

Note that since all t-reads of T_1 are executed sequentially, all these RAW/AWAR patterns are pairwise non-overlapping, which completes the proof. □

4.6 Related work and Discussion

In this section, we summarize the complexity bounds for blocking TMs presented in this chapter and identify some open questions.

Sequential TMs. Theorem 4.2 improves the read-validation step-complexity lower bound [62, 64] derived for *strict-data partitioning* (a very strong version of DAP) and *invisible reads*. In a *strict data partitioned* TM, the set of base objects used by the TM is split into disjoint sets, each storing information only about a single data item. Indeed, every TM implementation that is strict data-partitioned satisfies weak DAP, but not vice-versa (cf. Section 2.6). The definition of invisible reads assumed in [62, 64] requires that a t-read operation does not apply nontrivial events in any execution. Theorem 4.2 however, assumes *weak* invisible reads, stipulating that t-read operations of a transaction T do not apply nontrivial events only when T is not concurrent with any other transaction. We believe that the TM-progress and TM-liveness restrictions as well as the definitions of DAP and invisible reads we consider for this result are the weakest possible assumptions that may be made. To the best of our knowledge, these assumptions cover every TM implementation that is subject to the validation step-complexity [36, 39, 79].

Progressive TMs. We summarize the known complexity bounds for progressive (and resp. strongly progressive) TMs in Table 4.1 (and resp. Table 4.2). Some questions remain open. Can the tight bounds on step complexity for progressive TMs in Corollaries 4.14 and 4.16 be extended to strongly progressive TMs?

Guerraoui and Kapalka [64] proved that it is impossible to implement strictly serializable strongly progressive TMs that provide *wait-free* TM-liveness (every t-operation returns a matching response within

TM-correctness	TM-liveness	DAP	Invisible reads	Read-write	Complexity
Opacity	ICF	weak	yes	yes	$\Theta(\|Rset\|^2)$ step-complexity
Strict serializability	ICF	weak	yes	yes	$\Theta(\|Rset\|)$ step-complexity for tryCommit
Opacity	WF	strict	yes	yes	$O(1)$ RAW/AWAR, $O(1)$ stalls for t-reads
Opacity	starvation-free	strict			$\Theta(\|Wset\|)$ protected data

Table 4.1: Complexity bounds for progressive TMs.

TM-correctness	TM-liveness	Invisible reads	rmw primitives	Complexity
Strict serializability	WF		read-write	Impossible
Strict serializability			read-write, conditional	$\Omega(n \log n)$ RMRs
Opacity	starvation-free	yes	read-write	$O(1)$ RAW/AWAR

Table 4.2: Complexity bounds for strongly progressive TMs.

a finite number of steps) using only read and write primitives. Algorithm 4.5 describes one means to circumvent this impossibility result by describing an opaque strongly progressive TM implementation from read-write base objects that provides starvation-free TM-liveness.

We conjecture that the lower bound of Theorem 4.21 on the RMR complexity is tight. Proving this remains an interesting open question.

Permissive TMs. Crain et al. [34] proved that a permissive opaque TM implementation cannot maintain invisible reads, which inspired the derivation of our lower bound on RAW/AWAR complexity in Section 4.5. Furthermore, [34] described a permissive VWC TM implementation that ensures that t-read operations do not perform nontrivial primitives, but the tryCommit invoked by a read-only transaction perform a linear (in the size of the transaction's data set) number of RAW/AWARs. Thus, an open question is whether there exists a linear lower bound on RAW/AWAR complexity for weaker (than opacity) TM-correctness properties of VWC and TMS1.

5

Complexity bounds for non-blocking TMs

5.1 Overview

In the previous chapter, we presented complexity bounds for *lock-based* blocking TMs. Early TM implementations such as the popular *DSTM* [79] however avoid using locks and provide non-blocking TM-progress. In this chapter, we present several complexity bounds for non-blocking TMs exemplified by *obstruction-freedom*, possibly the weakest non-blocking progress condition [78, 82].

We first establish that it is impossible to implement a strictly serializable obstruction-free TM that provides both weak DAP and *read invisibility*. Indeed, popular obstruction-free TMs like *DSTM* [79] and *FSTM* [52] are weak DAP, but use *visible* reads for aborting pending writing transactions. Secondly, we show that a t-read operation in a n-process strictly serializable obstruction-free TM implementation may incur $\Omega(n)$ stalls. Specifically, we prove that every such TM implementation has a $(n-1)$-stall execution for an invoked t-read operation. Thirdly, we prove that any RW DAP opaque obstruction-free TM implementation has an execution in which a read-only transaction incurs $\Omega(n)$ non-overlapping *RAWs* or *AWARs*. Finally, we show that there exists a considerable complexity gap between blocking (*i.e.*, progressive) and non-blocking (*i.e.*, obstruction-free) TM implementations. We use the progressive opaque TM implementation LP described in Algorithm 4.1 (Chapter 4) to establish a linear separation in memory stall and RAW/AWAR complexity between blocking and non-blocking TMs.

Formally, let $\mathcal{OF}$ denote the class of TMs that provide OF TM-progress and OF TM-liveness.

Roadmap of Chapter 5. In Section 5.2, we show that no strictly serializable TM in $\mathcal{OF}$ can be weak DAP and have invisible reads. In Section 5.3, we determine stall complexity bounds for strictly serializable TMs in $\mathcal{OF}$, and in Section 5.4, we present a linear (in n) lower bound on the RAW/AWAR complexity for RW DAP opaque TMs in $\mathcal{OF}$. In Section 5.5, we describe two obstruction-free algorithms: a RW DAP opaque TM and a weak DAP (but not RW DAP) opaque TM. In Section 5.6, we present complexity gaps between blocking and non-blocking TM implementations. We conclude this chapter with a discussion on related work and open questions concerning obstruction-free TMs.

5.2 Impossibility of weak DAP and invisible reads

In this section, we prove that it is impossible to combine weak DAP and invisible reads for strictly serializable TMs in $\mathcal{OF}$.

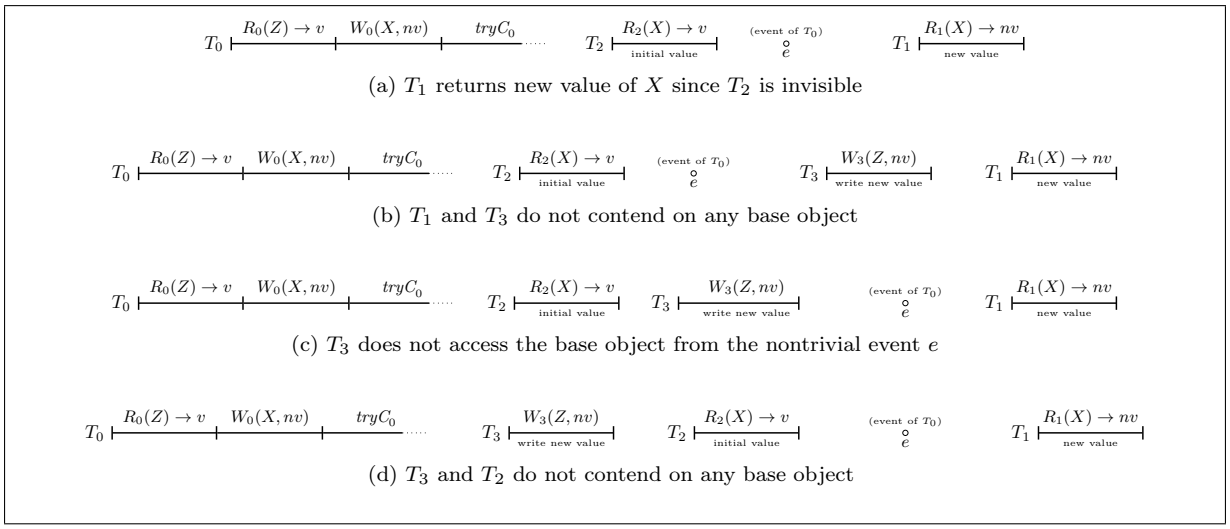
T_0 $R_0(Z) \to v$ $W_0(X, nv)$ $tryC_0$
T_2 $R_2(X) \to v$ initial value
(event of T_0) e
T_1 $R_1(X) \to nv$ new value
(a) T_1 returns new value of X since T_2 is invisible
T_0 $R_0(Z) \to v$ $W_0(X, nv)$ $tryC_0$
T_2 $R_2(X) \to v$ initial value
(event of T_0) e
T_3 $W_3(Z, nv)$ write new value
T_1 $R_1(X) \to nv$ new value
(b) T_1 and T_3 do not contend on any base object
T_0 $R_0(Z) \to v$ $W_0(X, nv)$ $tryC_0$
T_2 $R_2(X) \to v$ initial value
T_3 $W_3(Z, nv)$ write new value
(event of T_0) e
T_1 $R_1(X) \to nv$ new value
(c) T_3 does not access the base object from the nontrivial event e
T_0 $R_0(Z) \to v$ $W_0(X, nv)$ $tryC_0$
T_3 $W_3(Z, nv)$ write new value
T_2 $R_2(X) \to v$ initial value
(event of T_0) e
T_1 $R_1(X) \to nv$ new value
(d) T_3 and T_2 do not contend on any base object

Figure 5.1: Executions in the proof of Theorem 5.1; execution in 5.1d is not strictly serializable

Here is a proof sketch: suppose, by contradiction, that such a TM implementation M exists. Consider an execution E of M in which a transaction T_0 performs a t-read of t-object Z (returning the initial value v), writes nv (new value) to t-object X, and commits. Let E' denote the longest prefix of E that cannot be extended with the t-complete step contention-free execution of any transaction that reads nv in X and commits.

Thus if T_0 takes one more step, then the resulting execution $E' \cdot e$ can be extended with the t-complete step contention-free execution of a transaction T_1 that reads nv in X and commits.

Since M uses invisible reads, the following execution exists: E' can be extended with the t-complete step contention-free execution of a transaction T_2 that reads the initial value v in X and commits, followed by the step e of T_0 after which transaction T_1 running step contention-free reads nv in X and commits. Moreover, this execution is indistinguishable to T_1 and T_2 from an execution in which the read set of T_0 is empty. Thus, we can modify this execution by inserting the step contention-free execution of a committed transaction T_3 that writes a new value to Z after E', but preceding T_2 in real-time order. Intuitively, by weak DAP, transactions T_1 and T_2 cannot distinguish this execution from the original one in which T_3 does not participate.

Thus, we can show that the following execution exists: E' is extended with the t-complete step contention-free execution of T_3 that writes nv to Z and commits, followed by the t-complete step contention-free execution of T_2 that reads the initial value v in X and commits, followed by the step e of T_0, after which T_1 reads nv in X and commits.

This execution is, however, not strictly serializable: T_0 must appear in any serialization (T_1 reads a value written by T_0). Transaction T_2 must precede T_0, since the t-read of X by T_2 returns the initial value of X. To respect real-time order, T_3 must precede T_2. Finally, T_0 must precede T_3 since the t-read of Z returns the initial value of Z. The cycle $T_0 \to T_3 \to T_2 \to T_0$ implies a contradiction.

The formal proof follows.

Theorem 5.1. *There does not exist a weak DAP strictly serializable TM implementation in $\mathcal{OF}$ that uses invisible reads.*

Proof. By contradiction, assume that such an implementation $M \in \mathcal{OF}$ exists. Let v be the initial value of t-objects X and Z. Consider an execution E of M in which a transaction T_0 performs $read_0(Z) \to v$ (returning v), writes $nv \neq v$ to X, and commits. Let E' denote the longest prefix of E that cannot be extended with the t-complete step contention-free execution of any transaction performing a t-read X that returns nv and commits.

Let e be the enabled event of transaction T_0 in the configuration after E'. Without loss of generality, assume that $E' \cdot e$ can be extended with the t-complete step contention-free execution of a committed transaction T_1 that reads X and returns nv. Let $E' \cdot e \cdot E_1$ be such an execution, where E_1 is the t-complete step contention-free execution fragment of transaction T_1 that performs $read_1(X) \to nv$ and commits.

We now prove that M has an execution of the form $E' \cdot E_2 \cdot e \cdot E_1$, where E_2 is the t-complete step contention-free execution fragment of transaction T_2 that performs $read_2(X) \to v$ and commits.

We observe that $E' \cdot E_2$ is an execution of M. Indeed, by OF TM-progress and OF TM-liveness, T_2 must return a matching response that is not A_2 in $E' \cdot E_2$, and by the definition of E', this response must be the initial value v of X.

By the assumption of invisible reads, E_2 does not contain any nontrivial events. Consequently, $E' \cdot E_2 \cdot e \cdot E_1$ is indistinguishable to transaction T_1 from the execution $E' \cdot e \cdot E_1$. Thus, $E' \cdot E_2 \cdot e \cdot E_1$ is also an execution of M (Figure 5.1a).

Claim 5.2. *M has an execution of the form $E' \cdot E_2 \cdot E_3 \cdot e \cdot E_1$ where E_3 is the t-complete step contention-free execution fragment of transaction T_3 that writes $nv \neq v$ to Z and commits.*

Proof. The proof is through a sequence of indistinguishability arguments to construct the execution.

We first claim that M has an execution of the form $E' \cdot E_2 \cdot e \cdot E_1 \cdot E_3$. Indeed, by OF TM-progress and OF TM-liveness, T_3 must be committed in $E' \cdot E_2 \cdot e \cdot E_1 \cdot E_3$.

Since M uses invisible reads, the execution $E' \cdot E_2 \cdot e \cdot E_1 \cdot E_3$ is indistinguishable to transactions T_1 and T_3 from the execution $\hat{E} \cdot E_1 \cdot E_3$, where $\hat{E}$ is the t-incomplete step contention-free execution of transaction T_0 with $Wset_{\hat{E}}(T_0) = \{X\}$; $Rset_{\hat{E}}(T_0) = \emptyset$ that writes nv to X.

Observe that the execution $E' \cdot E_2 \cdot e \cdot E_1 \cdot E_3$ is indistinguishable to transactions T_1 and T_3 from the execution $\hat{E} \cdot E_1 \cdot E_3$, in which transactions T_3 and T_1 are disjoint-access. Consequently, by Lemma 2.10, T_1 and T_3 do not contend on any base object in $\hat{E} \cdot E_1 \cdot E_3$. Thus, M has an execution of the form $E' \cdot E_2 \cdot e \cdot E_3 \cdot E_1$ (Figure 5.1b).

By definition of E', T_0 applies a nontrivial primitive to some base object, say b, in event e that T_1 must access in E_1. Thus, the execution fragment E_3 does not contain any nontrivial event on b in the execution $E' \cdot E_2 \cdot e \cdot E_1 \cdot E_3$. In fact, since T_3 is disjoint-access with T_0 in the execution $\hat{E} \cdot E_3 \cdot E_1$, by Lemma 2.10, it cannot access the base object b to which T_0 applies a nontrivial primitive in the event e. Thus, transaction T_3 must perform the same sequence of events E_3 immediately after E', implying that M has an execution of the form $E' \cdot E_2 \cdot E_3 \cdot e \cdot E_1$ (Figure 5.1c). □

Finally, we observe that the execution $E' \cdot E_2 \cdot E_3 \cdot e \cdot E_1$ established in Claim 5.2 is indistinguishable to transactions T_2 and T_3 from an execution $\tilde{E} \cdot E_2 \cdot E_3 \cdot e \cdot E_1$, where $Wset(T_0) = \{X\}$ and $Rset(T_0) = \emptyset$ in $\tilde{E}$. But transactions T_3 and T_2 are disjoint-access in $\tilde{E} \cdot E_2 \cdot E_3 \cdot e \cdot E_1$ and by Lemma 2.10, T_2 and T_3 do not contend on any base object in this execution. Thus, M has an execution of the form $E' \cdot E_3 \cdot E_2 \cdot e \cdot E_1$ (Figure 5.1d) in which T_3 precedes T_2 in real-time order.

However, the execution $E' \cdot E_3 \cdot E_2 \cdot e \cdot E_1$ is not strictly serializable: T_0 must be committed in any serialization and transaction T_2 must precede T_0 since $read_2(X)$ returns the initial value of X. To respect real-time order, T_3 must precede T_2, while T_0 must precede T_1 since $read_1(X)$ returns nv, the value of X updated by T_0. Finally, T_0 must precede T_3 since $read_0(Z)$ returns the initial value of Z. But there exists no such serialization—contradiction. □

5.3 A linear lower bound on memory stall complexity

We prove a linear (in n) lower bound for strictly serializable TM implementations in $\mathcal{OF}$ on the total number of *memory stalls* incurred by a single t-read operation.

Inductively, for each $k \leq n-1$, we construct a specific *k-stall execution* [46] in which some t-read operation by a process p incurs k stalls. In the k-stall execution, k processes are partitioned into disjoint subsets $S_1, \ldots, S_i$. The execution can be represented as $\alpha \cdot \sigma_1 \cdots \sigma_i$; α is p-free, where in each σ_j, $j = 1, \ldots, i$, p first runs by itself, then each process in S_j applies a *nontrivial* event on a base object b_j, and then p applies an event on b_j. Moreover, p does not detect step contention in this execution and, thus, must return a non-abort value in its t-read and commit in the solo extension of it. Additionally, it is guaranteed that in any extension of α by the processes other than $\{p\} \cup S_1 \cup S_2 \cup \ldots \cup S_i$, no nontrivial primitive is applied on a base object accessed in $\sigma_1 \cdots \sigma_i$.

Assuming that $k \leq n-2$, we introduce a not previously used process executing an updating transaction immediately after α, so that the subsequent t-read operation executed by p is "perturbed" (must return another value). This will help us to construct a $(k+k')$-stall execution $\alpha \cdot \alpha' \cdot \sigma_1 \cdots \sigma_i \cdot \sigma_{i+1}$, where $k' > 0$.

The formal proof follows:

Theorem 5.3. *Every strictly serializable TM implementation $M \in \mathcal{OF}$ has a $(n-1)$-stall execution E for a t-read operation performed in E.*

Proof. We proceed by induction. Observe that the empty execution is a 0-stall execution since it vacuously satisfies the invariants of Definition 2.18.

Let v be the initial value of t-objects X and Z. Let $\alpha = \alpha_1 \cdots \alpha_{n-2}$ be a step contention-free execution of a strictly serializable TM implementation $M \in \mathcal{OF}$, where for all $j \in \{1, \ldots, n-2\}$, α_j is the longest prefix of the execution fragment $\bar{\alpha}_j$ that denotes the t-complete step-contention free execution of committed transaction T_j (invoked by process p_j) that performs $read_j(Z) \to v$, writes value $nv \neq v$ to X in the execution $\alpha_1 \cdots \alpha_{j-1} \cdot \bar{\alpha}_j$ such that

- $tryC_j()$ is incomplete in α_j,
- $\alpha_1 \cdots \alpha_j$ cannot be extended with the t-complete step contention-free execution fragment of any transaction T_{n-1} or T_n that performs exactly one t-read of X that returns nv and commits.

Assume, inductively, that $\alpha \cdot \sigma_1 \cdots \sigma_i$ is a k-stall execution for $read_n(X)$ executed by process p_n, where $0 \leq k \leq n-2$. By Definition 2.18, there are distinct base objects $b_1, \ldots b_i$ accessed by disjoint sets of processes $S_1 \ldots S_i$ in the execution fragment $\sigma_1 \cdots \sigma_i$, where $|S_1 \cup \ldots \cup S_i| = k$ and $\sigma_1 \cdots \sigma_i$ contains no events of processes not in $S_1 \cup \ldots \cup S_i \cup \{p_n\}$. We will prove that there exists a $(k+k')$-stall execution for $read_n(X)$, for some $k' \geq 1$.

By Lemma 2.12, $\alpha \cdot \sigma_1 \cdots \sigma_i$ is indistinguishable to T_n from a step contention-free execution. Let σ be the finite step contention-free execution fragment that extends $\alpha \cdot \sigma_1 \cdots \sigma_i$ in which T_n performs events by itself: completes $read_n(X)$ and returns a response. By OF TM-progress and OF TM-liveness, $read_n(X)$ and the subsequent $tryC_k$ must each return non-A_n responses in $\alpha \cdot \sigma_1 \cdots \sigma_i \cdot \sigma$. By construction of α and strict serializability of M, $read_n(X)$ must return the response v or nv in this execution. We prove that there exists an execution fragment γ performed by some process $p_{n-1} \notin (\{p_n\} \cup S_1 \cup \cdots \cup S_i)$ extending α that contains a nontrivial event on some base object that must be accessed by $read_n(X)$ in $\sigma_1 \cdots \sigma_i \cdot \sigma$.

Consider the case that $read_n(X)$ returns the response nv in $\alpha \cdot \sigma_1 \cdots \sigma_i \cdot \sigma$. We define a step contention-free fragment γ extending α that is the t-complete step contention-free execution of transaction T_{n-1} executed by some process $p_{n-1} \notin (\{p_n\} \cup S_1 \cup \cdots \cup S_i)$ that performs $read_{n-1}(X) \to v$, writes $nv \neq v$ to Z and commits. By definition of α, OF TM-progress and OF TM-liveness, M has an execution of the form $\alpha \cdot \gamma$. We claim that the execution fragment γ must contain a nontrivial event on some base object that must be accessed by $read_n(X)$ in $\sigma_1 \cdots \sigma_i \cdot \sigma$. Suppose otherwise. Then, $read_n(X)$ must return the response nv in $\sigma_1 \cdots \sigma_i \cdot \sigma$. But the execution $\alpha \cdot \sigma_1 \cdots \sigma_i \cdot \sigma$ is not strictly serializable. Since

$read_n(X) \to nv$, there exists a transaction $T_q \in txns(\alpha)$ that must be committed and must precede T_n in any serialization. Transaction T_{n-1} must precede T_n in any serialization to respect the real-time order and T_{n-1} must precede T_q in any serialization. Also, T_q must precede T_{n-1} in any serialization. But there exists no such serialization.

Consider the case that $read_n(X)$ returns the response v in $\alpha \cdot \sigma_1 \cdots \sigma_i \cdot \sigma$. In this case, we define the step contention-free fragment γ extending α as the t-complete step contention-free execution of transaction T_{n-1} executed by some process $p_{n-1} \notin (\{p_n\} \cup S_1 \cup \cdots \cup S_i)$ that writes $nv \neq v$ to X and commits. By definition of α, OF TM-progress and OF TM-liveness, M has an execution of the form $\alpha \cdot \gamma$. By strict serializability of M, the execution fragment γ must contain a nontrivial event on some base object that must be accessed by $read_n(X)$ in $\sigma_1 \cdots \sigma_i \cdot \sigma$. Suppose otherwise. Then, $\sigma_1 \cdots \sigma_i \cdot \gamma \cdot \sigma$ is an execution of M in which $read_n(X) \to v$. But this execution is not strictly serializable: every transaction $T_q \in txns(\alpha)$ must be aborted or must be preceded by T_n in any serialization, but committed transaction T_{n-1} must precede T_n in any serialization to respect the real-time ordering of transactions. But then $read_n(X)$ must return the new value nv of X that is updated by T_{n-1}—contradiction.

Since, by Definition 2.18, the execution fragment γ executed by some process $p_{n-1} \notin (\{p_n\} \cup S_1 \cup \cdots \cup S_i)$ contains no nontrivial events to any base object accessed in $\sigma_1 \cdots \sigma_i$, it must contain a nontrivial event to some base object $b_{i+1} \notin \{b_1, \ldots, b_i\}$ that is accessed by T_n in the execution fragment σ.

Let $\mathcal{A}$ denote the set of all finite $(\{p_n\} \cup S_1 \ldots \cup S_i)$-free execution fragments that extend α. Let $b_{i+1} \notin \{b_1, \ldots, b_i\}$ be the first base object accessed by T_n in the execution fragment σ to which some transaction applies a nontrivial event in the execution fragment $\alpha' \in \mathcal{A}$. Clearly, some such execution $\alpha \cdot \alpha'$ exists that contains a nontrivial event in α' to some distinct base object b_{i+1} not accessed in the execution fragment $\sigma_1 \cdots \sigma_i$. We choose the execution $\alpha \cdot \alpha' \in \mathcal{A}$ that maximizes the number of transactions that are poised to apply nontrivial events on b_{i+1} in the configuration after $\alpha \cdot \alpha'$. Let S_{i+1} denote the set of processes executing these transactions and $k' = |S_{i+1}|$ ($k' > 0$ as already proved).

We now construct a $(k + k')$-stall execution $\alpha \cdot \alpha' \cdot \sigma_1 \cdots \sigma_i \cdot \sigma_{i+1}$ for $read_n(X)$, where in σ_{i+1}, p_n applies events by itself, then each of the processes in S_{i+1} applies a nontrivial event on b_{i+1}, and finally, p_n accesses b_{i+1}.

By construction, $\alpha \cdot \alpha'$ is p_n-free. Let σ_{i+1} be the prefix of σ not including T_n's first access to b_{i+1}, concatenated with the nontrivial events on b_{i+1} by each of the k' transactions executed by processes in S_{i+1} followed by the access of b_{i+1} by T_n. Observe that T_n performs exactly one t-operation $read_n(X)$ in the execution fragment $\sigma_1 \cdots \sigma_{i+1}$ and $\sigma_1 \cdots \sigma_{i+1}$ contains no events of processes not in $(\{p_n\} \cup S_1 \cup \cdots \cup S_i \cup S_{i+1})$.

To complete the induction, we need to show that in every $(\{p_n\} \cup S_1 \cup \cdots \cup S_i \cup S_{i+1})$-free extension of $\alpha \cdot \alpha'$, no transaction applies a nontrivial event to any base object accessed in the execution fragment $\sigma_1 \cdots \sigma_i \cdot \sigma_{i+1}$. Let β be any such execution fragment that extends $\alpha \cdot \alpha'$. By our construction, σ_{i+1} is the execution fragment that consists of events by p_n on base objects accessed in $\sigma_1 \cdots \sigma_i$, nontrivial events on b_{i+1} by transactions in S_{i+1} and finally, an access to b_{i+1} by p_n. Since $\alpha \cdot \sigma_1 \cdots \sigma_i$ is a k-stall execution by our induction hypothesis, $\alpha' \cdot \beta$ is $(\{p_n\} \cup S_1 \ldots \cup S_i\})$-free and thus, $\alpha' \cdot \beta$ does not contain nontrivial events on any base object accessed in $\sigma_1 \cdots \sigma_i$. We now claim that β does not contain nontrivial events to b_{i+1}. Suppose otherwise. Thus, there exists some transaction T' that has an enabled nontrivial event to b_{i+1} in the configuration after $\alpha \cdot \alpha' \cdot \beta'$, where β' is some prefix of β. But this contradicts the choice of $\alpha \cdot \alpha'$ as the extension of α that maximizes k'.

Thus, $\alpha \cdot \alpha' \cdot \sigma_1 \cdots \sigma_i \cdot \sigma_{i+1}$ is indeed a $(k+k')$-stall execution for T_n where $1 < k < (k+k') \leq (n-1)$. □

Since there are at most n processes that are concurrent at any prefix of an execution, the lower bound of Theorem 5.3 is tight.

5.4 A linear lower bound on expensive synchronization for RW DAP

We prove that opaque, RW DAP TM implementations in $\mathcal{OF}$ have executions in which some read-only transaction performs a linear (in n) number of non-overlapping RAWs or AWARs.

Prior to presenting the formal proof, we present an overview (the executions used in the proof are depicted in Figure 5.2).

We first construct an execution of the form $\bar{\rho}_1 \cdots \bar{\rho}_m$, where for all $j \in \{1, \ldots, m\}$; $m = n - 3$, $\bar{\rho}_j$ denotes the t-complete step contention-free execution of transaction T_j that reads the initial value v in a distinct t-object Z_j, writes a new value nv to a distinct t-object X_j and commits. Observe that since any two transactions that participate in this execution are mutually read-write disjoint-access, they cannot contend on the same base object and, thus, the execution appears solo to each of them.

Let each of two new transactions T_{n-1} and T_n perform m t-reads on objects $X_1, \ldots, X_m$. For $j \in \{1, \ldots, m\}$, we now define ρ_j to be the longest prefix of $\bar{\rho}_j$ such that $\rho_1 \cdots \rho_j$ cannot be extended the complete step contention-free execution fragment of T_{n-1} or T_n where the t-read of X_j returns nv (Figure 5.2a). Let e_j be the event by T_j enabled after $\rho_1 \cdots \rho_j$. Let us count the number of indices $j \in \{1, \ldots, m\}$ such that T_{n-1} (resp., T_n) reads the new value nv in X_j when it runs after $\rho_1 \cdots \rho_j \cdot e_j$. Without loss of generality, assume that T_{n-1} has more such indices j than T_n. We are going to show that, in the worst-case, T_n must perform $\lceil \frac{m}{2} \rceil$ non-overlapping RAW/AWARs in the course of performing m t-reads of $X_1, \ldots, X_m$ immediately after $\rho_1 \cdots \rho_m$.

Consider any $j \in \{1, \ldots, m\}$ such that T_{n-1}, when it runs step contention-free after $\rho_1 \cdots \rho_j \cdot e_j$, reads nv in X_j. We claim that, in $\rho_1 \cdots \rho_m$ extended with the step contention-free execution of T_n performing j t-reads $\mathit{read}_n(X_1) \cdots \mathit{read}_n(X_j)$, the t-read of X_j must contain a RAW or an AWAR.

Suppose not. Then we are going to schedule a specific execution of T_j and T_{n-1} concurrently with $\mathit{read}_n(X_j)$ so that T_n cannot detect the concurrency. By the definition of ρ_j and the fact that the TM is RW DAP, T_n, when it runs step contention-free after $\rho_1 \cdots \rho_m$, must read v (the initial value) in X_j (Figure 5.2b). Then the following execution exists: $\rho_1 \cdots \rho_m$ is extended with the t-complete step contention-free execution of T_{n-2} writing nv to Z_j and committing, after which T_n runs step contention-free and reads v in X_j (Figure 5.2c). Since, by the assumption, $\mathit{read}_n(X_j)$ contains no RAWs or AWARs, we show that we can run T_{n-1} performing j t-reads concurrently with the execution of $\mathit{read}_n(X_j)$ so that T_n and T_{n-1} are unaware of step contention and $\mathit{read}_{n-1}(X_j)$ still reads the value nv in X_j.

To understand why this is possible, consider the following: we take the execution depicted in Figure 5.2c, but without the execution of $\mathit{read}_n(X_j)$, *i.e*, $\rho_1 \cdots \rho_m$ is extended with the step contention-free execution of committed transaction T_{n-2} writing nv to Z_j, after which T_n runs step contention-free performing $j-1$ t-reads. This execution can be extended with the step e_j by T_j, followed by the step contention-free execution of transaction T_{n-1} in which it reads nv in X_j. Indeed, by RW DAP and the definition of $\rho_j \cdot e_j$, there exists such an execution (Figure 5.2d).

Since $\mathit{read}_n(X_j)$ contains no RAWs or AWARs, we can reschedule the execution fragment e_j followed by the execution of T_{n-1} so that it is concurrent with the execution of $\mathit{read}_n(X_j)$ and neither T_n nor T_{n-1} see a difference (Figure 5.2e). Therefore, in this execution, $\mathit{read}_n(X_j)$ still returns v, while $\mathit{read}_{n-1}(X_j)$ returns nv.

However, the resulting execution (Figure 5.2e) is not opaque. In any serialization the following must hold. Since T_{n-1} reads the value written by T_j in X_j, T_j must be committed. Since $\mathit{read}_n(X_j)$ returns the initial value v, T_n must precede T_j. The committed transaction T_{n-2}, which writes a new value to Z_j, must precede T_n to respect the real-time order on transactions. However, T_j must precede T_{n-2} since $\mathit{read}_j(Z_j)$ returns the initial value and the implementation is opaque. The cycle $T_j \to T_{n-2} \to T_n \to T_j$ implies a contradiction.

Thus, we can show that transaction T_n must perform $\Omega(n)$ RAW/AWARs during the execution of m t-reads immediately after $\rho_1 \cdots \rho_m$.

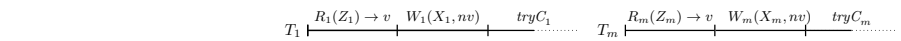

(a) Transactions in $\{T_1, \ldots, T_m\}$;$m = n - 3$ are mutually read-write disjoint-access and concurrent; they are poised to apply a nontrivial primitive

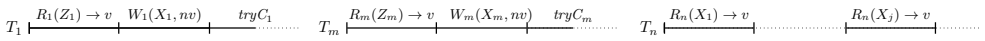

(b) T_n performs m reads; each $read_n(X_j)$ returns initial value v

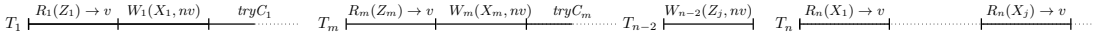

(c) T_{n-2} commits; T_n is read-write disjoint-access with T_{n-2}

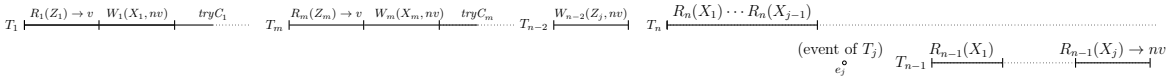

(d) T_{n-1} is read-write disjoint-access with T_{n-2}; $read_{n-1}(X_j)$ returns the value nv

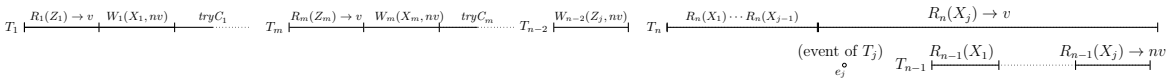
T_1 $R_1(Z_1) \to v$ $W_1(X_1, nv)$ $tryC_1$
T_m $R_m(Z_m) \to v$ $W_m(X_m, nv)$ $tryC_m$
T_{n-2} $W_{n-2}(Z_j, nv)$
T_n $R_n(X_1) \cdots R_n(X_{j-1})$
$R_n(X_j) \to v$
(event of T_j) e_j
T_{n-1} $R_{n-1}(X_1)$
$R_{n-1}(X_j) \to nv$

(e) Suppose $read_n(X_j)$ does not perform a RAW/AWAR, T_n and T_{n-1} are unaware of step contention and T_n misses the event of T_j, but $R_{n-1}(X_j)$ returns the value of X_j that is updated by T_j

Figure 5.2: Executions in the proof of Theorem 5.4; execution in 5.2e is not opaque

Theorem 5.4. *Every RW DAP opaque TM implementation $M \in \mathcal{OF}$ has an execution E in which some read-only transaction $T \in txns(E)$ performs $\Omega(n)$ non-overlapping RAW/AWARs.*

Proof. For all $j \in \{1, \ldots, m\}$; $m = n - 3$, let v be the initial value of t-objects X_j and Z_j. Throughout this proof, we assume that, for all $i \in \{1, \ldots, n\}$, transaction T_i is invoked by process p_i.

By OF TM-progress and OF TM-liveness, any opaque and RW DAP TM implementation $M \in \mathcal{OF}$ has an execution of the form $\bar{\rho}_1 \cdots \bar{\rho}_m$, where for all $j \in \{1, \ldots, m\}$, $\bar{\rho}_j$ denotes the t-complete step contention-free execution of transaction T_j that performs $read_j(Z_j) \to v$, writes value $nv \neq v$ to X_j and commits.

By construction, any two transactions that participate in $\bar{\rho}_1 \cdots \bar{\rho}_n$ are mutually read-write disjoint-access and cannot contend on the same base object. It follows that for all $1 \leq j \leq m$, $\bar{\rho}_j$ is an execution of M.

For all $j \in \{1, \ldots, m\}$, we iteratively define an execution ρ_j of M as follows: it is the longest prefix of $\bar{\rho}_j$ such that $\rho_1 \cdots \rho_j$ cannot be extended with the complete step contention-free execution fragment of transaction T_n that performs j t-reads: $read_n(X_1) \cdots read_n(X_j)$ in which $read_n(X_j) \to nv$ nor with the complete step contention-free execution fragment of transaction T_{n-1} that performs j t-reads: $read_{n-1}(X_1) \cdots read_{n-1}(X_j)$ in which $read_{n-1}(X_j) \to nv$ (Figure 5.2a).

For any $j \in \{1, \ldots, m\}$, let e_j be the event transaction T_j is poised to apply in the configuration after $\rho_1 \cdots \rho_j$. Thus, the execution $\rho_1 \cdots \rho_j \cdot e_j$ can be extended with the complete step contention-free executions of at least one of transaction T_n or T_{n-1} that performs j t-reads of $X_1, \ldots, X_j$ in which the t-read of X_j returns the new value nv. Let T_{n-1} be the transaction that must return the new value for the maximum number of X_j's when $\rho_1 \cdots \rho_j \cdot e_j$ is extended with the t-reads of $X_1, \ldots, X_j$. We show that, in the worst-case, transaction T_n must perform $\lceil \frac{m}{2} \rceil$ non-overlapping RAW/AWARs in the course of performing m t-reads of $X_1, \ldots, X_m$ immediately after $\rho_1 \cdots \rho_m$. Symmetric arguments apply for the case when T_n must return the new value for the maximum number of X_j's when $\rho_1 \cdots \rho_j \cdot e_j$ is extended with the t-reads of $X_1, \ldots, X_j$.

Proving the RAW/AWAR lower bound. We prove that transaction T_n must perform $\lceil \frac{m}{2} \rceil$ non-overlapping RAWs or AWARs in the course of performing m t-reads of $X_1, \ldots, X_m$ immediately after the execution $\rho_1 \cdots \rho_m$. Specifically, we prove that T_n must perform a RAW or an AWAR during the execution of the t-read of each X_j such that $\rho_1 \cdots \rho_j \cdot e_j$ can be extended with the complete step contention-free execution of T_{n-1} as it performs j t-reads of $X_1 \ldots X_j$ in which the t-read of X_j returns the new value nv. Let $\mathbb{J}$ denote the of all $j \in \{1, \ldots, m\}$ such that $\rho_1 \cdots \rho_j \cdot e_j$ extended with the complete step contention-free execution of T_{n-1} performing j t-reads of $X_1 \ldots X_j$ must return the new value nv during the t-read of X_j.

We first prove that, for all $j \in \mathbb{J}$, M has an execution of the form $\rho_1 \cdots \rho_m \cdot \delta_j$ (Figures 5.2a and 5.2b), where δ_j is the complete step contention-free execution fragment of T_n that performs j t-reads: $read_n(X_1) \cdots read_n(X_j)$, each of which return the initial value v.

By definition of ρ_j, OF TM-progress and OF TM-liveness, M has an execution of the form $\rho_1 \cdots \rho_j \cdot \delta_j$. By construction, transaction T_n is read-write disjoint-access with each transaction $T \in \{T_{j+1}, \ldots, T_m\}$ in $\rho_1 \cdots \rho_j \cdots \rho_m \cdot \delta_j$. Thus, T_n cannot contend with any of the transactions in $\{T_{j+1}, \ldots, T_m\}$, implying that, for all $j \in \{1, \ldots, m\}$, M has an execution of the form $\rho_1 \cdots \rho_m \cdot \delta_j$ (Figure 5.2b).

We claim that, for each $j \in \mathbb{J}$, the t-read of X_j performed by T_n must perform a RAW or an AWAR in the course of performing j t-reads of $X_1, \ldots, X_j$ immediately after $\rho_1 \cdots \rho_m$. Suppose by contradiction that $read_n(X_j)$ does not perform a RAW or an AWAR in $\rho_1 \cdots \rho_m \cdot \delta_m$.

Claim 5.5. *For all $j \in \mathbb{J}$, M has an execution of the form $\rho_1 \cdots \rho_j \cdots \rho_m \cdot \delta_{j-1} \cdot e_j \cdot \beta$ where, β is the complete step contention-free execution fragment of transaction T_{n-1} that performs j t-reads: $read_{n-1}(X_1) \cdots read_{n-1}(X_{j-1}) \cdot read_{n-1}(X_j)$ in which $read_{n-1}(X_j)$ returns nv.*

Proof. We observe that transaction T_n is read-write disjoint-access with every transaction $T \in \{T_j, T_{j+1}, \ldots, T_m\}$ in $\rho_1 \cdots \rho_j \cdots \rho_m \cdot \delta_{j-1}$. By RW DAP, it follows that M has an execution of the form $\rho_1 \cdots \rho_j \cdots \rho_m \cdot \delta_{j-1} \cdot e_j$ since T_n cannot perform a nontrivial event on the base object accessed by T_j in the event e_j.

By the definition of ρ_j, transaction T_{n-1} must access the base object to which T_j applies a nontrivial primitive in e_j to return the value nv of X_j as it performs j t-reads of $X_1, \ldots, X_j$ immediately after the execution $\rho_1 \cdots \rho_j \cdots \rho_m \cdot \delta_{j-1} \cdot e_j$. Thus, M has an execution of the form $\rho_1 \cdots \rho_j \cdot \delta_{j-1} \cdot e_j \cdot \beta$.

By construction, transactions T_{n-1} is read-write disjoint-access with every transaction $T \in \{T_{j+1}, \ldots, T_m\}$ in $\rho_1 \cdots \rho_j \cdots \rho_m \cdot \delta_{j-1} \cdot e_j \cdot \beta$. It follows that M has an execution of the form $\rho_1 \cdots \rho_j \cdots \rho_m \cdot \delta_{j-1} \cdot e_j \cdot \beta$. □

Claim 5.6. *For all $j \in \{1, \ldots, m\}$, M has an execution of the form $\rho_1 \cdots \rho_j \cdots \rho_m \cdot \gamma \cdot \delta_{j-1} \cdot e_j \cdot \beta$, where γ is the t-complete step contention-free execution fragment of transaction T_{n-2} that writes $nv \neq v$ to Z_j and commits.*

Proof. Observe that T_{n-2} precedes transactions T_n and T_{n-1} in real-time order in the above execution.

By OF TM-progress and OF TM-liveness, transaction T_{n-2} must be committed in $\rho_1 \cdots \rho_j \cdots \rho_m \cdot \gamma$.

Since transaction T_{n-1} is read-write disjoint-access with T_{n-2} in $\rho_1 \cdots \rho_j \cdots \rho_m \cdot \gamma \cdot \delta_{j-1} \cdot e_j \cdot \beta$, T_{n-1} does not contend with T_{n-2} on any base object (recall that we associate an edge with t-objects in the conflict graph only if they are both contained in the write set of some transaction). Since the execution fragment β contains an access to the base object to which T_j performs a nontrivial primitive in the event e_j, T_{n-2} cannot perform a nontrivial event on this base object in γ. It follows that M has an execution of the form $\rho_1 \cdots \rho_j \cdots \rho_m \cdot \gamma \cdot \delta_{j-1} \cdot e_j \cdot \beta$ since, it is indistinguishable to T_{n-1} from the execution $\rho_1 \cdots \rho_j \cdots \rho_m \cdot \delta_{j-1} \cdot e_j \cdot \beta$ (the existence of which is already established in Claim 5.5). □

Recall that transaction T_n is read-write disjoint-access with T_{n-2} in $\rho_1 \cdots \rho_j \cdots \rho_m \cdot \gamma \cdot \delta_j$. Thus, M has an execution of the form $\rho_1 \cdots \rho_j \cdots \rho_m \cdot \gamma \cdot \delta_j$ (Figure 5.2c).

Deriving a contradiction. For all $j \in \{1, \ldots, m\}$, we represent the execution fragment δ_j as $\delta_{j-1} \cdot \pi^j$, where π^j is the complete execution fragment of the j^{th} t-read $read_n(X_j) \to v$. By our assumption, π^j does not contain a RAW or an AWAR.

For succinctness, let $\alpha = \rho_1 \cdots \rho_m \cdot \gamma \cdot \delta_{j-1}$. We now prove that if π^j does not contain a RAW or an AWAR, we can define $\pi_1^j \cdot \pi_2^j = \pi^j$ to construct an execution of the form $\alpha \cdot \pi_1^j \cdot e_j \cdot \beta \cdot \pi_2^j$ (Figure 5.2e) such that

- no event in π_1^j is the application of a nontrivial primitive
- $\alpha \cdot \pi_1^j \cdot e_j \cdot \beta \cdot \pi_2^j$ is indistinguishable to T_n from the step contention-free execution $\alpha \cdot \pi_1^j \cdot \pi_2^j$
- $\alpha \cdot \pi_1^j \cdot e_j \cdot \beta \cdot \pi_2^j$ is indistinguishable to T_{n-1} from the step contention-free execution $\alpha \cdot e_j \cdot \beta$.

The following claim defines π_1^j and π_2^j to construct this execution.

Claim 5.7. *For all $j \in \{1, \ldots, m\}$, M has an execution of the form $\alpha \cdot \pi_1^j \cdot e_j \cdot \beta \cdot \pi_2^j$.*

Proof. Let t be the first event containing a write to a base object in the execution fragment π^j. We represent π^j as the execution fragment $\pi_1^j \cdot t \cdot \pi_f^j$. Since π_1^j does not contain nontrivial events that write to a base object, $\alpha \cdot \pi_1^j \cdot e_j \cdot \beta$ is indistinguishable to transaction T_{n-1} from the step contention-free execution $\alpha \cdot e_j \cdot \beta$ (as already proven in Claim 5.6). Consequently, $\alpha \cdot \pi_1^j \cdot e_j \cdot \beta$ is an execution of M.

Since t is not an atomic-write-after-read, M has an execution of the form $\alpha \cdot \gamma \cdot \pi_1^j \cdot e_j \cdot \beta \cdot t$. Secondly, since π^j does not contain a read-after-write, any read of a base object performed in π_f^j may only be performed to base objects previously written in $t \cdot \pi_f^j$. Thus, $\alpha \cdot \pi_1^j \cdot e_j \cdot \beta \cdot t \cdot \pi_f^j$ is indistinguishable to T_n from the step contention-free execution $\alpha \cdot \pi_1^j \cdot t \cdot \pi_f^j$. But, as already proved, $\alpha \cdot \pi^j$ is an execution of M.

Choosing $\pi_2^j = t \cdot \pi_f^j$, it follows that M has an execution of the form $\alpha \cdot \pi_1^j \cdot e_j \cdot \beta \cdot \pi_2^j$. □

We have now proved that, for all $j \in \{1, \ldots, m\}$, M has an execution of the form $\rho_1 \cdots \rho_m \cdot \gamma \cdot \delta_{j-1} \cdot \pi_1^j \cdot e_j \cdot \beta \cdot \pi_2^j$ (Figure 5.2e).

The execution in Figure 5.2e is not opaque. Indeed, in any serialization the following must hold. Since T_{n-1} reads the value written by T_j in X_j, T_j must be committed. Since $read_n(X_j)$ returns the initial value v, T_n must precede T_j. The committed transaction T_{n-2}, which writes a new value to Z_j, must precede T_n to respect the real-time order on transactions. However, T_j must precede T_{n-2} since $read_j(Z_j)$ returns the initial value of Z_j. The cycle $T_j \to T_{n-2} \to T_n \to T_j$ implies that there exists no such a serialization.

Thus, for each $j \in \mathbb{J}$, transaction T_n must perform a RAW or an AWAR during the t-read of X_j in the course of performing m t-reads of $X_1, \ldots, X_m$ immediately after $\rho_1 \cdots \rho_m$. Since $|\mathbb{J}| \geq \lceil \frac{(n-3)}{2} \rceil$, in the worst-case, T_n must perform $\Omega(n)$ RAW/AWARs during the execution of m t-reads immediately after $\rho_1 \cdots \rho_m$. □

5.5 Algorithms for obstruction-free TMs

In this section, we present two opaque obstruction-free TM implementations: the first one satisfies RW DAP, but not strict DAP while the second one satisfies weak DAP, but not RW DAP.

5.5.1 An opaque RW DAP TM implementation

In this section, we describe a RW DAP TM implementation in $\mathcal{OF}$ (based on *DSTM* [79]).

Every t-object X_m maintains a base object $tvar[m]$ and every transaction T_k maintains a $status[k]$ base object. Both base objects support the *read*, *write* and *compare-and-swap* (*cas*) primitives.

The object $tvar[m]$ stores a triple: the *owner* of X_m is an updating transaction that performs the latest write to X_m, the *old value* and *new value* of X_m represent two latest versions of X_m. The base object $status[k]$ denotes if T_k is *live* (i.e. t-incomplete), *committed* or *aborted*. Intuitively, if $status[k]$ is *committed*, then other transactions can safely read the value of the t-objects updated by T_k.

Implementation of $read_k(X_m)$ first reads $tvar[m]$ and checks if the *owner* of X_m is *live*; if so, it forcefully aborts the owning transaction and returns the *old value* of X_m. Otherwise, if the owner is *committed*, it returns the *new value* of X_m. In both cases, it only returns a non-abort value if no t-object previously read has been updated since. The $write_k(X_m, v)$ works similar to the $read_k(X_m)$ implementation; but additionally, if the *owner* of X_m is *live*, it forcefully aborts the owning transaction, assumes ownership of X_m, sets v as the new value of X_m and leaves the *old value* of X_m unchanged. Otherwise, if the *owner* of X_m is a committed transaction, it updates the *old value* of X_m to be the value of X_m updated by its previous *owner*. The $tryC_k$ implementation sets $status[k]$ to *committed* if it has not been set to aborted by a concurrent transaction, otherwise T_k is deemed aborted. Since any t-read operation performs at most two AWARs and the *tryC* performs only a single AWAR, any read-only transaction T performs at most $O(|Rset(T)|)$ AWARs. The pseudocode is described in Algorithm 5.1.

Lemma 5.8. *Algorithm 5.1 implements an opaque TM.*

Proof. Since opacity is a safety property, we only consider finite executions [18]. Let E by any finite execution of Algorithm 5.1. Let $<_E$ denote a total-order on events in E.

Let H denote a subsequence of E constructed by selecting *linearization points* of t-operations performed in E. The linearization point of a t-operation *op*, denoted as ℓ_{op} is associated with a base object event or an event performed during the execution of *op* using the following procedure.

Completions. First, we obtain a completion of E by removing some pending invocations and adding responses to the remaining pending invocations involving a transaction T_k as follows: every incomplete $read_k$, $write_k$, $tryC_k$ operation is removed from E; an incomplete $write_k$ is removed from E.

Linearization points. We now associate linearization points to t-operations in the obtained completion of E as follows:

- For every t-read op_k that returns a non-A_k value, ℓ_{op_k} is chosen as the event in Line 12 of Algorithm 5.1, else, ℓ_{op_k} is chosen as invocation event of op_k
- For every t-write op_k that returns a non-A_k value, ℓ_{op_k} is chosen as the event in Line 37 of Algorithm 5.1, else, ℓ_{op_k} is chosen as invocation event of op_k
- For every $op_k = tryC_k$ that returns C_k, ℓ_{op_k} is associated with Line 65.

$<_H$ denotes a total-order on t-operations in the complete sequential history H.

Serialization points. The serialization of a transaction T_j, denoted as δ_{T_j} is associated with the linearization point of a t-operation performed during the execution of the transaction.

We obtain a t-complete history $\bar{H}$ from H as follows: for every transaction T_k in H that is complete, but not t-complete, we insert $tryC_k \cdot A_k$ after H.

$\bar{H}$ is thus a t-complete sequential history. A t-complete t-sequential history S equivalent to $\bar{H}$ is obtained by associating serialization points to transactions in $\bar{H}$ as follows:

- If T_k is an update transaction that commits, then δ_{T_k} is ℓ_{tryC_k}

Algorithm 5.1 RW DAP opaque implementation $M \in \mathcal{OF}$; code for T_k

```
1:  Shared base objects:
2:    tvar[m], storing [owner_m, oval_m, nval_m]
3:        for each t-object X_m, supports read, write, cas
4:        owner_m, a transaction identifier
5:        oval_m ∈ V
6:        nval_m ∈ V
7:    status[k] ∈ {live, aborted, committed},
8:        for each T_k; supports read, write, cas
9:  Local variables:
10:   Rset_k, Wset_k for every transaction T_k;
11:       dictionaries storing {X_m, Tvar[m]}

12: read_k(X_m):
13:   [owner_m, oval_m, nval_m] ← tvar[m].read()
14:   if owner_m ≠ k then
15:     s_m ← status[owner_m].read()
16:     if s_m = committed then
17:       curr = nval_m
18:     else if s_m = aborted then
19:       curr = oval_m
20:     else
21:       if status[owner_m].cas(live, aborted) then
22:         curr = oval_m
23:       else
24:         Return A_k
25:     if status[k] = live ∧ ¬validate() then
26:       Rset(T_k).add({X_m, [owner_m, oval_m, nval_m]})
27:       Return curr
28:     Return A_k
29:   else
30:     Return Rset(T_k).locate(X_m)

31: Function: validate():
32:   if ∃{X_j, [owner_j, oval_j, nval_j]} ∈ Rset(T_k):
33:       ([owner_j, oval_j, nval_j] ≠ tvar[j].read()) then
34:     Return true
35:   Return false

36: write_k(X_m, v):
37:   [owner_m, oval_m, nval_m] ← tvar[m].read()
38:   if owner_m ≠ k then
39:     s_m ← status[owner_m].read()
40:     if s_m = committed then
41:       curr = nval_m
42:     else if s_m = aborted then
43:       curr = oval_m
44:     else
45:       if status[owner_m].cas(live, aborted) then
46:         curr = oval_m
47:       else
48:         Return A_k
49:     o_m ← tvar[m].cas([owner_m, oval_m, nval_m], [k, curr, v])

50:     if o_m ∧ status[k] = live then
51:       Wset_k.add({X_m, [k, curr, v]})
52:       Return ok
53:     else
54:       Return A_k
55:   else
56:     [owner_m, oval_m, nval_m] = Wset_k.locate(X_m)
57:     s = tvar[m].cas([owner_m, oval_m, nval_m], [k, oval_m, v])
58:     if s then
59:       Wset(T_k).add({X_m, [k, oval_m, v]})
60:       Return ok
61:     else
62:       Return A_k

63: tryC_k():
64:   if validate() then
65:     Return A_k
66:   if status[k].cas(live, committed) then
67:     Return C_k
68:   Return A_k
```

- If T_k is an aborted or read-only transaction in $\bar{H}$, then δ_{T_k} is assigned to the linearization point of the last t-read that returned a non-A_k value in T_k

$<_S$ denotes a total-order on transactions in the t-sequential history S.

Claim 5.9. *If $T_i \prec_H^{RT} T_j$, then $T_i <_S T_j$.*

Proof. This follows from the fact that for a given transaction, its serialization point is chosen between the first and last event of the transaction implying if $T_i \prec_H T_j$, then $\delta_{T_i} <_E \delta_{T_j}$ implies $T_i <_S T_j$ □

Claim 5.10. *If transaction T_i returns C_i in E, then* status[i]=committed *in E.*

Proof. Transaction T_i must perform the event in Line 66 before returning T_i i.e. the *cas* on its own *status* to change the value to *committed*. The proof now follows from the fact that any other transaction may change the *status* of T_i only if it is *live* (Lines 45 and 21). □

Claim 5.11. *S is legal.*

Proof. Observe that for every $read_j(X) \to v$, there exists some transaction T_i that performs $write_i(X, v)$ and completes the event in Line 22 to write v as the *new value* of X such that $read_j(X) \not\prec_H^{RT} write_i(X, v)$. For any updating committing transaction T_i, $\delta_{T_i} = \ell_{tryC_i}$. Since $read_j(X)$ returns a response v, the event

in Line 12 must succeed the event in Line 66 when T_i changes *status[i]* to *committed*. Suppose otherwise, then $read_j(X)$ subsequently forces T_i to abort by writing *aborted* to *status[i]* and must return the *old value* of X that is updated by the previous *owner* of X, which must be committed in E (Line 40). Since $\delta_{T_i} = \ell_{tryC_i}$ precedes the event in Line 66, it follows that $\delta_{T_i} <_E \ell_{read_j(X)}$.

We now need to prove that $\delta_{T_i} <_E \delta_{T_j}$. Consider the following cases:

- if T_j is an updating committed transaction, then δ_{T_j} is assigned to ℓ_{tryC_j}. But since $\ell_{read_j(X)} <_E \ell_{tryC_j}$, it follows that $\delta_{T_i} <_E \delta_{T_j}$.
- if T_j is a read-only or aborted transaction, then δ_{T_j} is assigned to the last t-read that did not abort. Again, it follows that $\delta_{T_i} <_E \delta_{T_j}$.

To prove that S is legal, we need to show that, there does not exist any transaction T_k that returns C_k in S and performs $write_k(X, v')$; $v' \neq v$ such that $T_i <_S T_k <_S T_j$. Now, suppose by contradiction that there exists a committed transaction T_k, $X \in Wset(T_k)$ that writes $v' \neq v$ to X such that $T_i <_S T_k <_S T_j$. Since T_i and T_k are both updating transactions that commit,

$$(T_i <_S T_k) \iff (\delta_{T_i} <_E \delta_{T_k})$$
$$(\delta_{T_i} <_E \delta_{T_k}) \iff (\ell_{tryC_i} <_E \ell_{tryC_k})$$

Since, T_j reads the value of X written by T_i, one of the following is true: $\ell_{tryC_i} <_E \ell_{tryC_k} <_E \ell_{read_j(X)}$ or $\ell_{tryC_i} <_E \ell_{read_j(X)} <_E \ell_{tryC_k}$.

If $\ell_{tryC_i} <_E \ell_{tryC_k} <_E \ell_{read_j(X)}$, then the event in Line 66 performed by T_k when it changes the status field to *committed* precedes the event in Line 12 performed by T_j. Since $\ell_{tryC_i} <_E \ell_{tryC_k}$ and both T_i and T_k are committed in E, T_k must perform the event in Line 37 after T_i changes *status[i]* to *committed* since otherwise, T_k would perform the event in Line 45 and change *status[i]* to *aborted*, thereby forcing T_i to return A_i. However, $read_j(X)$ observes that the *owner* of X is T_k and since the *status* of T_k is committed at this point in the execution, $read_j(X)$ must return v' and not v—contradiction.

Thus, $\ell_{tryC_i} <_E \ell_{read_j(X)} <_E \ell_{tryC_k}$. We now need to prove that δ_{T_j} indeed precedes $\delta_{T_k} = \ell_{tryC_k}$ in E.

Now consider two cases:

- Suppose that T_j is a read-only transaction. Then, δ_{T_j} is assigned to the last t-read performed by T_j that returns a non-A$_j$ value. If $read_j(X)$ is not the last t-read that returned a non-A$_j$ value, then there exists a $read_j(X')$ such that $\ell_{read_j(X)} <_E \ell_{tryC_k} <_E \ell_{read_j(X')}$. But then this t-read of X' must abort since the value of X has been updated by T_k since T_j first read X—contradiction.
- Suppose that T_j is an updating transaction that commits, then $\delta_{T_j} = \ell_{tryC_j}$ which implies that $\ell_{read_j(X)} <_E \ell_{tryC_k} <_E \ell_{tryC_j}$. Then, T_j must neccesarily perform the validation of its read set in Line 65 and return A_j—contradiction.

□

Claims 5.9 and 5.11 establish that Algorithm 5.1 is opaque. □

Theorem 5.12. *Algorithm 5.1 describes a RW DAP, progressive opaque TM implementation $M \in \mathcal{OF}$ such that in every execution E of M,*

- *the total number of stalls incurred by a t-read operation invoked in E is $O(n)$,*
- *every read-only transaction $T \in txns(E)$ performs $O(|Rset(T)|)$ AWARs in E, and*
- *every complete t-read operation invoked by transaction $T_k \in txns(E)$ performs $O(|Rset_E(T_k)|$ steps.*

Proof. (Opacity) Follows from Lemma 5.8

(TM-liveness and TM-progress) Since none of the implementations of the t-operations in Algorithm 5.1 contain unbounded loops or waiting statements, every t-operation op_k returns a matching response after taking a finite number of steps. Thus, Algorithm 5.1 provides wait-free TM-liveness.

To prove OF TM-progress, we proceed by enumerating the cases under which a transaction T_k may be aborted in any execution.

- Suppose that there exists a $read_k(X_m)$ performed by T_k that returns A_k. If $read_k(X_m)$ returns A_k in Line 28, then there exists a concurrent transaction that updated a t-object in $Rset(T_k)$ or changed *status[k]* to *aborted*. In both cases, T_k returns A_k only because there is step contention.
- Suppose that there exists a $write_k(X_m, v)$ performed by T_k that returns A_k in Line 54. Thus, either a concurrent transaction has changed *status[k]* to *aborted* or the value in $tvar[m]$ has been updated since the event in Line 37. In both cases, T_k returns A_k only because of step contention with another transaction.
- Suppose that a $read_k(X_m)$ or $write_k(X_m, v)$ return A_k in Lines 21 and 45 respectively. Thus, a concurrent transaction has takes steps concurrently by updating the *status* of $owner_m$ since the read by T_k in Lines 12 and 37 respectively.
- Suppose that $tryC_k()$ returns A_k in Line 62. This is because there exists a t-object in $Rset(T_k)$ that has been updated by a concurrent transaction since, *i.e.*, $tryC_k()$ returns A_k only on encountering step contention.

It follows that in any step contention-free execution of a transaction T_k from a T_k-free execution, T_k must return C_k after taking a finite number of steps.

The enumeration above also proves that M implements a progressive TM.

(Read-write disjoint-access parallelism) Consider any execution E of Algorithm 5.1 and let T_i and T_j be any two transactions that contend on a base object b in E. We need to prove that there is a path between a t-object in $Dset(T_i)$ and a t-object in $Dset(T_j)$ in $\tilde{G}(T_i, T_j, E)$ or there exists $X \in Dset(T_i) \cap Dset(T_j)$. Recall that there exists an edge between t-objects X and Y in $\tilde{G}(T_i, T_j, E)$ only if there exists a transaction $T \in txns(E)$ such that $\{X, Y\} \in Wset(T)$.

- Suppose that T_i and T_j contend on base object *tvar*[m] belonging to t-object X_m in E. By Algorithm 5.1, a transaction accesses X_m only if X_m is contained in $Dset(T_m)$. Thus, both T_i and T_j must access X_m.
- Suppose that T_i and T_j contend on base object *status*[i] in E (the case when T_i and T_j contend on *status*[j] is symmetric). T_j accesses *status[i]* while performing a t-read of some t-object X in Lines 15 and 21 only if T_i is the *owner* of X. Also, T_j accesses *status[i]* while performing a t-write to X in Lines 39 and 45 only if T_i is the *owner* of X. But if T_i is the *owner* of X, then $X \in Wset(T_i)$.
- Suppose that T_i and T_j contend on base object *status*[m] belonging to some transaction T_m in E. Firstly, observe that T_i or T_j access *status[m]* only if there exist t-objects X and Y in $Dset(T_i)$ and $Dset(T_j)$ respectively such that $\{X, Y\} \in Wset(T_m)$. This is because T_i and T_j would both read *status[m]* in Lines 15 (during t-read) and 39 (during t-write) only if T_m was the previous *owner* of X and Y. Secondly, one of T_i or T_j applies a nontrivial primitive to *status[m]* only if T_i and T_j read *status[m]=live* in Lines 15 (during t-read) and 37 (during t-write). Thus, at least one of T_i or T_j is concurrent to T_m in E. It follows that there exists a path between X and Y in $\tilde{G}(T_i, T_j, E)$.

(Complexity) Every t-read operation performs at most one AWAR in an execution E (Line 21) of Algorithm 5.1. It follows that any read-only transaction $T_k \in txns(E)$ performs at most $|Rset(T_k)|$ AWARs in E.

The linear step-complexity is immediate from the fact that during the t-read operations, the transaction validates its entire read set (Line 25). All other t-operations incur $O(1)$ step-complexity since they involve no iteration statements like *for* and *while* loops.

Since at most $n-1$ transactions may be t-incomplete at any point in an execution E, it follows that E is at most a $(n-1)$-stall execution for any t-read *op* and every $T \in txns(E)$ incurs $O(n)$ stalls on account of any event performed in E. More specifically, consider the following execution E: for all $i \in \{1, \ldots, n-1\}$, each transaction T_i performs $write_i(X_m, v)$ in a step-contention free execution until it is poised to apply a nontrivial event on $tvar[m]$ (Line 22). By OF TM-progress, we construct E such that each of the T_i is poised to apply a nontrivial event on $tvar[m]$ after E. Consider the execution fragment of $read_n(X_m)$ that is poised to perform an event e that reads $tvar[m]$ (Line 12) immediately after E. In the constructed execution, T_n incurs $O(n)$ stalls on account of e and thus, produces the desired $(n-1)$-stall execution for $read_n(X)$. □

5.5.2 An opaque weak DAP TM implementation

In this section, we describe a weak DAP TM implementation in $\mathcal{OF}$ with constant step-complexity t-read operations.

Algorithm 5.2 Weak DAP opaque implementation $M \in \mathcal{OF}$; code for T_k

```
1:  read_k(X_m):
2:    [owner_m, oval_m, nval_m] ← tvar[m].read()
3:    if owner_m ≠ k then
4:      s_m ← status[owner_m].read()
5:      if s_m = committed then
6:        curr = nval_m
7:      else if s_m = aborted then
8:        curr = oval_m
9:      else
10:       if status[owner_m].cas(live, aborted) then
11:         curr = oval_m
12:       Return A_k
13:     o_m ← tvar[m].cas([owner_m, oval_m, nval_m], [k, oval_m, nval_m])
14:     if o_m ∧ status[k] = live then
15:       Rset(T_k).add({X_m, [owner_m, oval_m, nval_m]})
16:       Return curr
17:   else
18:     Return Rset(T_k).locate(X_m)

19: tryC_k():
20:   if status[k].cas(live, committed) then
21:     Return C_k
22:   Return A_k
```

Algorithm 5.2 describes a weak DAP implementation in $\mathcal{OF}$ that does not satisfy read-write DAP. The code for the t-write operations is identical to Algorithm 5.1. During the t-read of t-object X_m by transaction T_k, T_k becomes the *owner* of X_m thus eliminating the per-read validation step-complexity inherent to Algorithm 5.1. Similarly, $tryC_k$ also not involve performing the validation of the T_k's read set; the implementation simply sets $status[k] = committed$ and returns C_k.

Theorem 5.13. *Algorithm 5.2 describes a weak TM implementation $M \in \mathcal{OF}$ such that in any execution E of M, for every transaction $T \in txns(E)$, T performs $O(1)$ steps during the execution of any t-operation in E.*

Proof. The proofs of opacity, TM-liveness and TM-progress are almost identical to the analogous proofs for Algorithm 5.1.

(Weak disjoint-access parallelism) Consider any execution E of Algorithm 5.2 and let T_i and T_j be any two transactions that contend on a base object b in E. We need to prove that there is a path between a t-object in $Dset(T_i)$ and a t-object in $Dset(T_j)$ in $\tilde{G}(T_i, T_j, E)$ or there exists $X \in Dset(T_i) \cap Dset(T_j)$. Recall that there exists an edge between t-objects X and Y in $G(T_i, T_j, E)$ only if there exists a transaction $T \in txns(E)$ such that $\{X, Y\} \in Dset(T)$.

- Suppose that T_i and T_j contend on base object *tvar*[m] belonging to t-object X_m in E. By Algorithm 5.2, a transaction accesses X_m only if X_m is contained in $Dset(T_m)$. Thus, both T_i and T_j must access X_m.
- Suppose that T_i and T_j contend on base object *status*[i] in E (the case when T_i and T_j contend on *status*[j] is symmetric). T_j accesses *status[i]* while performing a t-read of some t-object X in Lines 4 and 10 only if T_i is the *owner* of X. Also, T_j accesses *status[i]* while performing a t-write to X in Lines 39 and 45 only if T_i is the *owner* of X. But if T_i is the *owner* of X, then $X \in Dset(T_i)$.
- Suppose that T_i and T_j contend on base object *status*[m] belonging to some transaction T_m in E. Firstly, observe that T_i or T_j access *status[m]* only if there exist t-objects X and Y in $Dset(T_i)$ and $Dset(T_j)$ respectively such that $\{X, Y\} \in Dset(T_m)$. This is because T_i and T_j would both read *status[m]* in Lines 4 (during t-read) and 39 (during t-write) only if T_m was the previous *owner* of X and Y. Secondly, one of T_i or T_j applies a nontrivial primitive to *status[m]* only if T_i and T_j read *status[m]=live* in Lines 4 (during t-read) and 37 (during t-write). Thus, at least one of T_i or T_j is concurrent to T_m in E. It follows that there exists a path between X and Y in $\tilde{G}(T_i, T_j, E)$.

(Complexity) Since no implementation of any of the t-operation contains any iteration statements like *for* and *while* loops), the proof follows. □

5.6 Why Transactional memory should not be obstruction-free

	Obstruction-free TMs	Progressive TM LP
strict DAP	No [60]	Yes
invisible reads+weak DAP	No	Yes
stall complexity of t-reads	$\Omega(n)$	$O(1)$
RAW/AWAR complexity	$\Omega(n)$	$O(1)$
read-write base objects, wait-free TM-liveness	No [64]	Yes

Figure 5.3: Complexity gap between blocking and non-blocking TMs

As a synchronization abstraction, TM came as an alternative to conventional lock-based synchronization, and it therefore appears natural that early TM implementations [52, 79, 101, 120], avoided using locks. Instead, early TM designs relied on non-blocking synchronization, where a prematurely halted transaction cannot prevent all other transactions from committing. Possibly the weakest progress condition elucidating non-blocking TM-progress is obstruction-freedom.

However, in 2005, Ennals [48] argued that obstruction-free TMs inherently yield poor performance, because they require transactions to forcefully abort each other. Ennals further described a *lock-based* TM implementation [47] satisfying progressiveness that he claimed to outperform *DSTM* [79], the most referenced obstruction-free TM implementation at the time. Inspired by [48], more recent lock-based progressive TMs, such as *TL* [40], *TL2* [39] and *NOrec* [36], demonstrate better performance than obstruction-free TMs on most workloads.

There is a considerable amount of empirical evidence on the performance gap between non-blocking (obstruction-free) and blocking (progressive) TM implementations but no analytical result explains it. We present complexity lower and upper bounds that provide such an explanation.

To exhibit a complexity gap between blocking and non-blocking TMs, we go back to the the progressive opaque TM implementation LP (Algorithm 4.1) that beats the impossibility result and the lower bounds we established for obstruction-free TMs. Recall that our implementation LP, (1) uses only read-write base objects and provides wait-free TM-liveness, (2) ensures strict DAP, (3) has invisible reads, (4) performs $O(1)$ non-overlapping RAWs/AWARs per transaction, and (5) incurs $O(1)$ memory stalls for read operations (Theorem 4.13). In contrast, from prior work and our lower bounds we know that (i) no OF TM that provides wait-free transactional operations can be implemented using only read-write base objects [64]; (ii) no OF TM can provide strict DAP [60]; (iii) no weak DAP OF TM has invisible reads

(Section 5.2) and (iv) no OF TM ensures a constant number of stalls incurred by a t-read operation (Section 5.3). Finally, (v) no RW DAP *opaque* OF TM has constant RAW/AWAR complexity (Section 5.4). In fact, (iv) and (v) exhibit a linear separation between blocking and non-blocking TMs w.r.t expensive synchronization and memory stall complexity, respectively.

Altogether, our results exhibit a considerable complexity gap between progressive and obstruction-free TMs, as summarized in Figure 5.3, that seems to justify the shift in TM practice (circa. 2005) from non-blocking to blocking TMs.

Overcoming our lower bounds for obstruction-free TMs individually is comparatively easy. Say, TL [40] combines strict DAP with invisible reads, but it is not read-write, and it does not provide constant RAW/AWAR and stall complexities.

Coming out with a single algorithm that beats all these lower bounds is quite nontrivial. Our algorithm *LP* incurs the cost of *incremental validation*, *i.e.*, checking that the current read set has not changed per every new read operation. This is, however, unavoidable for invisible read algorithms (cf. Theorem 4.2), and is, in fact, believed to yield better performance in practice than "visible" reads [36, 40, 47], and we show that it enables constant stall and RAW/AWAR complexity.

5.7 Related work and Discussion

In this section, we summarize the results presented in this chapter and identify some unresolved questions.

Lower bounds for non-blocking TMs. Complexity of obstruction-free TMs was first studied by Guerraoui and Kapalka [60, 64] who proved that they cannot provide strict DAP. However, as we show in Section 5.5, it is possible to realize weaker than strict DAP variants of obstruction-free opaque TMs. Bushkov et al. [31] improved on the impossibility result in [60] and showed that a variant of strict DAP cannot be combined with obstruction-free TM-progress, even if a weaker (than strictly serializability) TM-correctness property is assumed. In the thesis, we do not consider relaxations of strict serializability.

Guerraoui and Kapalka [60, 64] also proved that a strict serializable TM that provides OF TM-progress and wait-free TM-liveness cannot be implemented using only read and write primitives. An interesting open question is whether we can implement strict serializable TMs in $\mathcal{OF}$ using only read and write primitives.

Observe that, since there are at most n concurrent transactions, we cannot do better than $(n-1)$ stalls (cf. Definition 2.18). Thus, the lower bound of Theorem 5.3 is tight.

Moreover, we conjecture that the linear (in n) lower bound of Theorem 5.4 for RW DAP opaque obstruction-free TMs can be strengthened to be linear in the size of the transaction's read set. Then, Algorithm 5.1, which proves a linear upper bound in the size of the transaction's read set, would allow us to establish a linear tight bound (in the size of the transaction's read set) for RW DAP opaque obstruction-free TMs.

Blocking versus non-blocking TMs. As highlighted in [40, 48], obstruction-free TMs typically must forcefully abort pending conflicting transactions. This observation inspires the impossibility of invisible reads (Theorem 5.1). Typically, to detect the presence of a conflicting transaction and abort it, the reading transaction must employ a RAW or a read-modify-write primitive like *compare-and-swap*, motivating the linear lower bound on expensive synchronization (Theorem 5.4). Also, in obstruction-free TMs, a transaction may not wait for a concurrent inactive transaction to complete and, as a result, we may have an execution in which a transaction incurs a distinct stall due to a transaction run by each other process, hence the linear stall complexity (Theorem 5.3). Intuitively, since transactions in progressive TMs may abort themselves in case of conflicts, they can employ invisible reads and maintain constant stall and RAW/AWAR complexities.

The lower bound and the proof technique in Theorem 5.3 is inspired by an analogous lower bound on *linearizable solo-terminating* implementations [16, 46] of a wide class of "perturbable" objects that

include *counters*, *compare-and-swap* and *single-writer snapshots* [16, 46]. Informally, the definition of solo-termination (adapted to the TM context) says that for every finite execution E, and every transaction T that is t-incomplete in E, there is a finite step contention-free extension in which T eventually commits. Observe that, under this definition, T is guaranteed to commit even in some executions that are not step contention-free for T. However, the definition of OF TM-progress used in the thesis ensures that T is guaranteed to commit only if all its events are issued in the absence of step contention. Moreover, [16] described a single-lock (only the process holding the lock can invoke an operation) implementation of these objects that incurs $O(\log n)$ stalls, thus establishing a separation between the worst-case operation stall complexity of non-blocking and blocking (*i.e.*, lock-based) implementations of these objects. In this chapter, we presented a linear separation in memory stall complexity between obstruction-free TMs and lock-based TMs characterized by progressiveness, which is a strictly stronger (than single-lock) progress guarantee, thus establishing the inherent cost of non-blocking progress in the TM context.

Some benefits of obstruction-free TMs, namely their ability to make progress even if some transactions prematurely fail, are not provided by progressive TMs. However, several papers [39, 40, 48] argued that lock-based TMs tend to outperform obstruction-free ones by allowing for simpler algorithms with lower overhead, and their inherent progress issues may be resolved using timeouts and *contention-managers* [115]. This chapter explains the empirically observed performance gap between blocking and non-blocking TMs via a series of lower bounds on obstruction-free TMs and a progressive TM algorithm that beats all of them.

6

Lower bounds for partially non-blocking TMs

6.1 Overview

It is easy to see that *dynamic* TMs where the patterns in which transactions access t-objects are not known in advance do not allow for *wait-free* TMs [64], *i.e.*, every transaction must commit in a finite number of steps of the process executing it, regardless of the behavior of concurrent processes. Suppose that a transaction T_1 reads t-object X, then a concurrent transaction T_2 reads t-object Y, writes to X and commits, and finally T_2 writes to Y. Since T_1 has read the "old" value in X and T_2 has read the "old" value in Y, there is no way to commit T_1 and order the two transactions in a sequential execution. As this scenario can be repeated arbitrarily often, even the weaker guarantee of *local progress* that only requires that each transaction *eventually* commits if repeated sufficiently often, cannot be ensured by *any* strictly serializable TM implementation, regardless of the base objects it uses [32].[1]

But can we ensure that at least *some* transactions commit wait-free and what are the inherent costs? It is often argued that many realistic workloads are *read-dominated*: the proportion of read-only transactions is higher than that of updating ones, or read-only transactions have much larger data sets than updating ones [25, 65]. Therefore, it seems natural to require that read-only transactions commit wait-free. Since we are interested in complexity lower bounds, we require that updating transaction provide only sequential TM-progress.

First, we focus on strictly serializable TMs with the above TM-progress conditions that use invisible reads. We show that this requirement results in maintaining unbounded sets of versions for every data item, *i.e.*, such implementations may not be practical due to their space complexity. Secondly, we prove that strictly serializable TMs with these progress conditions cannot ensure strict DAP. Thus, two transactions that access mutually disjoint data sets may prevent each other from committing. Thirdly, for weak DAP TMs, we show that a read-only transaction (with an arbitrarily large read set) must sometimes perform at least one expensive synchronization pattern [17] per t-read operation, *i.e.*, the expensive synchronization complexity of a read-only transaction is linear in the size of its data set.

Formally, we denote by $\mathcal{RWF}$ the class of partially non-blocking TMs.

Definition 6.1. *(The class $\mathcal{RWF}$) A TM implementation $M \in \mathcal{RWF}$* iff *in its every execution:*

[1] Note that the counter-example would not work if we imagine that the data sets accessed by a transaction can be known in advance. However, in the thesis, we consider the conventional dynamic TM programming model.

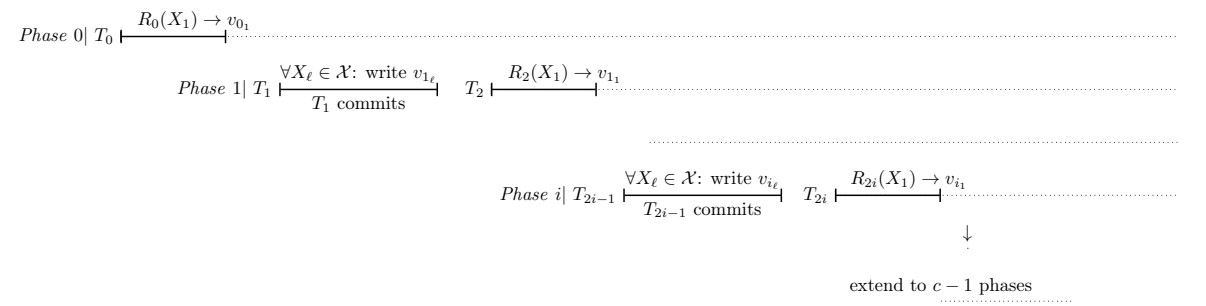
Phase 0| T_0 $R_0(X_1) \to v_{0_1}$
Phase 1| T_1 $\forall X_\ell \in \mathcal{X}$: write v_{1_ℓ}
T_1 commits
T_2 $R_2(X_1) \to v_{1_1}$
Phase i| T_{2i-1} $\forall X_\ell \in \mathcal{X}$: write v_{i_ℓ}
T_{2i-1} commits
T_{2i} $R_{2i}(X_1) \to v_{i_1}$
extend to $c-1$ phases

(a) for all $i \in \{1, \ldots, c-1\}$, T_{2i-1} writes v_{i_ℓ} to each X_ℓ; $read_{2i}(X_1)$ must return v_{i_1}

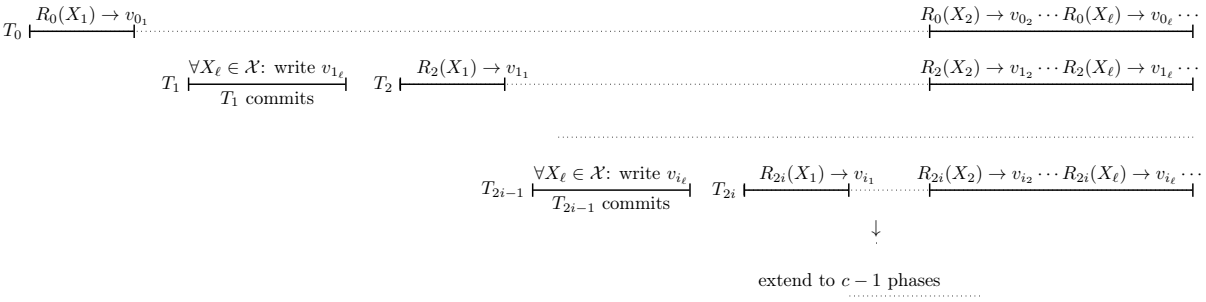
T_0 $R_0(X_1) \to v_{0_1}$
$R_0(X_2) \to v_{0_2} \cdots R_0(X_\ell) \to v_{0_\ell} \cdots$
T_1 $\forall X_\ell \in \mathcal{X}$: write v_{1_ℓ}
T_1 commits
T_2 $R_2(X_1) \to v_{1_1}$
$R_2(X_2) \to v_{1_2} \cdots R_2(X_\ell) \to v_{1_\ell} \cdots$
T_{2i-1} $\forall X_\ell \in \mathcal{X}$: write v_{i_ℓ}
T_{2i-1} commits
T_{2i} $R_{2i}(X_1) \to v_{i_1}$
$R_{2i}(X_2) \to v_{i_2} \cdots R_{2i}(X_\ell) \to v_{i_\ell} \cdots$
extend to $c-1$ phases

(b) extend every read-only transaction T_{2i} in phase i with t-reads of $X_2, \ldots X_\ell, \ldots$; each $read_{2i}(X_\ell)$ must return v_{i_ℓ}

Figure 6.1: Executions in the proof of Theorem 6.1; execution in 6.1a must maintain c distinct values of every t-object

- *(*wait-free TM-progress for read-only transactions*) every read-only transaction commits in a finite number of its steps, and*
- *(*sequential TM-progress and sequential TM-liveness for updating transactions*)* i.e.*, every transaction running step contention-free from a t-quiescent configuration, commits in a finite number of its steps.*

Roadmap of Chapter 6. Section 6.2 presents a lower bound on the inherent space complexity of TMs in $\mathcal{RWF}$. Section 6.3 proves the impossibility of strict DAP TMs in $\mathcal{RWF}$ while in Section 6.4, assuming weak DAP, we prove a linear, in the size of the transaction's read set, lower bound on expensive synchronization complexity. We conclude this chapter with a discussion of the related work and open questions concerning TMs in $\mathcal{RWF}$.

6.2 The space complexity of invisible reads

We prove that every strictly serializable TM implementation $M \in \mathcal{RWF}$ that uses invisible reads must keep unbounded sets of values for every t-object. To do so, for every $c \in \mathbb{N}$, we construct an execution of M that *maintains at least c distinct values for every t-object.* We require the following technical definition:

Definition 6.2. *Let E be any execution of a TM implementation M. We say that E* maintains c distinct values $\{v_1, \ldots, v_c\}$ of t-object X*, if there exists an execution $E \cdot E'$ of M such that*

- *E' contains the complete executions of c t-reads of X and,*
- *for all $i \in \{1, \ldots, c\}$, the response of the i^{th} t-read of X in E' is v_i.*

Theorem 6.1. *Let M be any strictly serializable TM implementation in $\mathcal{RWF}$ that uses invisible reads, and $\mathcal{X}$, any set of t-objects. Then, for every $c \in \mathbb{N}$, there exists an execution E of M such that E maintains at least c distinct values of each t-object $X \in \mathcal{X}$.*

Proof. Let v_{0_ℓ} be the initial value of t-object $X_\ell \in \mathcal{X}$. For every $c \in \mathbb{N}$, we iteratively construct an execution E of M of the form depicted in Figure 6.1a. The construction of E proceeds in phases: there are at most $c-1$ phases. For all $i \in \{0, \dots c-1\}$, we denote the execution after phase i as E_i which is defined as follows:

- E_0 is the complete step contention-free execution fragment α_0 of read-only transaction T_0 that performs $read_0(X_1) \to v_{0_1}$
- for all $i \in \{1, \dots, c-1\}$, E_i is defined to be an execution of the form $\alpha_0 \cdot \rho_1 \cdot \alpha_1 \cdots \rho_i \cdot \alpha_i$ such that for all $j \in \{1, \dots, i\}$,
  - ρ_j is the t-complete step contention-free execution fragment of an updating transaction T_{2j-1} that, for all $X_\ell \in \mathcal{X}$ writes the value v_{j_ℓ} and commits
  - α_j is the complete step contention-free execution fragment of a read-only transaction T_{2j} that performs $read_{2j}(X_1) \to v_{j_1}$

Since read-only transactions are invisible, for all $i \in \{0, \dots, c-1\}$, the execution fragment α_i does not contain any nontrivial events. Consequently, for all $i < j \leq c-1$, the configuration after E_i is indistinguishable to transaction T_{2j-1} from a t-quiescent configuration and it must be committed in ρ_j (by sequential progress for updating transactions). Observe that, for all $1 \leq j < i$, $T_{2j-1} \prec_E^{RT} T_{2i-1}$. Strict serializability of M now stipulates that, for all $i \in \{1, \dots, c-1\}$, the t-read of X_1 performed by transaction T_{2i} in the execution fragment α_i must return the value v_{i_1} of X_1 as written by transaction T_{2i-1} in the execution fragment ρ_i (in any serialization, T_{2i-1} is the latest committed transaction writing to X_1 that precedes T_{2i}). Thus, M indeed has an execution E of the form depicted in Figure 6.1a.

Consider the execution fragment E' that extends E in which, for all $i \in \{0, \dots, c-1\}$, read-only transaction T_{2i} is extended with the complete execution of the t-reads of every t-object $X_\ell \in \mathcal{X} \setminus \{X_1\}$ (depicted in Figure 6.1b).

We claim that, for all $i \in \{0, \dots, c-1\}$, and for all $X_\ell \in \mathcal{X} \setminus \{X_1\}$, $read_{2i}(X_\ell)$ performed by transaction T_{2i} must return the value v_{i_ℓ} of X_ℓ written by transaction T_{2i-1} in the execution fragment ρ_i. Indeed, by wait-free progress, $read_i(X_\ell)$ must return a non-abort response in such an extension of E. Suppose by contradiction that $read_i(X_\ell)$ returns a response that is not v_{i_ℓ}. There are two cases:

- $read_{2i}(X_\ell)$ returns the value v_{j_ℓ} written by transaction T_{2j-1}; $j < i$. However, since for all $j < i$, $T_{2j} \prec_E^{RT} T_{2i}$, the execution is not strictly serializable—contradiction.
- $read_{2i}(X_\ell)$ returns the value v_{j_ℓ} written by transaction T_{2j}; $j > i$. Since $read_i(X_1)$ returns the value v_{i_1} and $T_{2i} \prec_E^{RT} T_{2j}$, there exists no such serialization—contradiction.

Thus, E maintains at least c distinct values of every t-object $X \in \mathcal{X}$. □

6.3 Impossibility of strict DAP

In this section, we prove that it is impossible to derive strictly serializable TM implementations in $\mathcal{RWF}$ which ensure that any two transactions accessing pairwise disjoint data sets can execute without contending on the same base object.

Theorem 6.2. *There exists no strictly serializable strict DAP TM implementation in $\mathcal{RWF}$.*

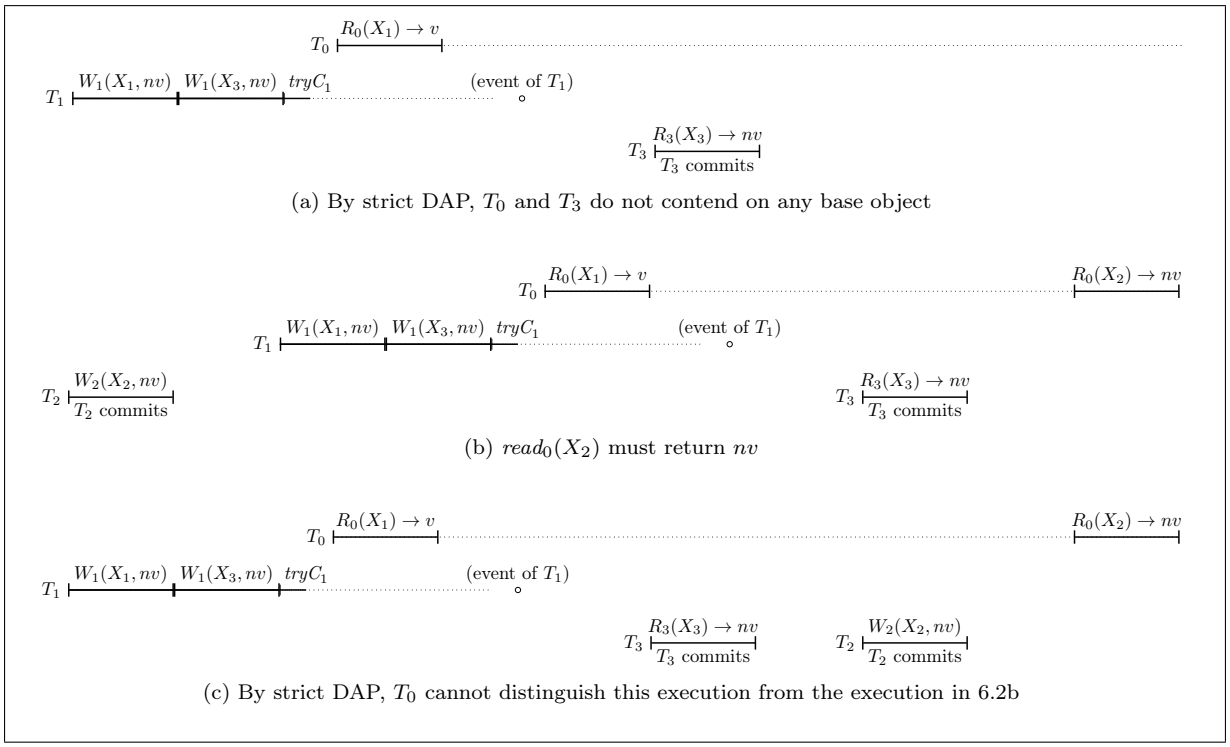

(a) By strict DAP, T_0 and T_3 do not contend on any base object

(b) $read_0(X_2)$ must return nv

(c) By strict DAP, T_0 cannot distinguish this execution from the execution in 6.2b

Figure 6.2: Executions in the proof of Theorem 6.2; execution in 6.2c is not strictly serializable

Proof. Suppose by contradiction that there exists a strict DAP TM implementation $M \in \mathcal{RWF}$.

Let v be the initial value of t-objects X_1, X_2 and X_3. Let π be the t-complete step contention-free execution of transaction T_1 that writes the value $nv \neq v$ to t-objects X_1 and X_3. By sequential progress for updating transactions, T_1 must be committed in π.

Note that any read-only transaction that runs step contention-free after some prefix of π must return a non-abort value. Since any such transaction reading X_1 or X_3 must return v after the empty prefix of π and nv when it starts from π, there exists π', the longest prefix of π that cannot be extended with the t-complete step contention-free execution of any transaction that performs a t-read of X_1 and returns nv nor with the t-complete step contention-free execution of any transaction that performs a t-read of X_3 and returns nv.

Consider the execution fragment $\pi' \cdot \alpha_1$, where α_1 is the complete step contention-free execution of transaction T_0 that performs $read_0(X_1) \to v$. Indeed, by definition of π' and wait-free progress (assumed for read-only transactions), M has an execution of the form $\pi' \cdot \alpha_1$.

Let e be the enabled event of transaction T_1 in the configuration after π'. Without loss of generality, assume that $\pi' \cdot e$ can be extended with the t-complete step contention-free execution of a transaction that reads X_3 and returns nv.

We now prove that M has an execution of the form $\pi' \cdot \alpha_1 \cdot e \cdot \beta \cdot \gamma$, where

- β is the t-complete step contention-free execution fragment of transaction T_3 that performs $read_3(X_3) \to nv$ and commits
- γ is the t-complete step contention-free execution fragment of transaction T_2 that writes nv to X_2 and commits.

Observe that, by definition of π', M has an execution of the form $\pi' \cdot e \cdot \beta$. By construction, transaction T_1 applies a nontrivial primitive to a base object, say b in the event e that is accessed by transaction T_3 in the execution fragment β. Since transactions T_0 and T_3 access mutually disjoint data sets in $\pi' \cdot \alpha_1 \cdot e \cdot \beta$, T_3 does not access any base object in β to which transaction T_0 applies a nontrivial primtive in the

execution fragment α_1 (assumption of strict DAP). Thus, α_1 does not contain a nontrivial primitive to b and $\pi' \cdot \alpha_1 \cdot e \cdot \beta$ is indistinguishable to T_3 from the execution $\pi' \cdot e \cdot \beta$. This proves that M has an execution of the form $\pi' \cdot \alpha_1 \cdot e \cdot \beta$ (depicted in Figure 6.2a).

Since transaction T_2 writes to t-object $Dset(T_2) = X_2 \notin \{Dset(T_1) \cup Dset(T_0) \cup Dset(T_3)\}$, by strict DAP, the configuration after $\pi' \cdot \alpha_1 \cdot e \cdot \beta$ is indistinguishable to T_2 from a t-quiescent configuration. Indeed, transaction T_2 does not contend with any of the transactions T_1, T_0 and T_3 on any base object in $\pi' \cdot \alpha_1 \cdot e \cdot \beta \cdot \gamma$. Sequential progress of M requires that T_2 must be committed in $\pi' \cdot \alpha_1 \cdot e \cdot \beta \cdot \gamma$. Thus, M has an execution of the form $\pi' \cdot \alpha_1 \cdot e \cdot \beta \cdot \gamma$.

By the above arguments, the execution $\pi' \cdot \alpha_1 \cdot e \cdot \beta \cdot \gamma$ is indistinguishable to each of the transactions T_1, T_0, T_2 and T_3 from $\gamma \cdot \pi' \cdot \alpha_1 \cdot e \cdot \beta$ in which transaction T_2 precedes T_1 in real-time ordering. Thus, $\gamma \cdot \pi' \cdot \alpha_1 \cdot e \cdot \beta$ is also an execution of M.

Consider the extension of the execution $\gamma \cdot \pi' \cdot \alpha_1 \cdot e \cdot \beta$ in which transaction T_0 performs $read_0(X_2)$ and commits (depicted in Figure 6.2b). Strict serializability of M stipulates that $read_0(X_2)$ must return nv since T_2 (which writes nv to X_2 in γ) precedes T_0 in this execution.

Similarly, we now extend the execution $\pi' \cdot \alpha_1 \cdot e \cdot \beta \cdot \gamma$ with the complete step contention-free execution fragment of the t-read of X_2 by transaction T_0. Since T_0 is a read-only transaction, it must be committed in this extension. However, as proved above, this execution is indistinguishable to T_0 from the execution depicted in Figure 6.2b in which $read_0(X_2)$ must return nv. Thus, M has an execution of the form $\pi' \cdot \alpha_1 \cdot e \cdot \beta \cdot \gamma \cdot \alpha_2$, where T_0 performs $read_0(X_2) \to nv$ in α_2 and commits.

However, the execution $\pi' \cdot \alpha_1 \cdot e \cdot \beta \cdot \gamma \cdot \alpha_2$ (depicted in Figure 6.2c) is not strictly serializable. Transaction T_1 must be committed in any serialization and must precede transaction T_3 since $read_3(X_3)$ returns the value of X_3 written by T_m. However, transaction T_0 must must precede T_1 since $read_0(X_1)$ returns the initial the value of X_1. Also, transaction T_2 must precede T_0 since $read_0(X_2)$ returns the value of X_2 written by T_2. But transaction T_3 must precede T_2 to respect real-time ordering of transactions. Thus, T_1 must precede T_0 in any serialization. But there exists no such serialization: a contradiction to the assumption that M is strictly serializable. □

6.4 A linear lower bound on expensive synchronization for weak DAP

In this section, we prove a linear lower bound (in the size of the transaction's read set) on the number of RAWs or AWARs for weak DAP TM implementations in $\mathcal{RWF}$. To do so, we construct an execution in which each t-read operation of an arbitrarily long read-only transaction contains a RAW or an AWAR.

Theorem 6.3. *Every strictly serializable weakly DAP TM implementation $M \in \mathcal{RWF}$ has, for all $m \in \mathbb{N}$, an execution in which some read-only transaction T_0 with $m = |Rset(T_0)|$ performs $\Omega(m)$ RAWs/AWARs.*

Proof. Let v be the initial value of each of the t-objects $X_1, \ldots, X_m$. Consider the t-complete step contention-free execution of transaction T_0 that performs m t-reads $read_0(X_1)$, $read_0(X_1)$,... $read_0(X_m)$ and commits. We prove that each of the first $m-1$ t-reads must perform a RAW or an AWAR.

For all $j \in \{1, \ldots, m-1\}$, M has an execution of the form $\alpha_1 \cdot \alpha_2 \cdots \alpha_j$, where for all $i \in \{1, \ldots, j\}$, α_i is the complete step contention-free execution fragment of $read_0(X_j) \to v$. Assume inductively that each of the first $j-1$ t-reads performs a RAW or an AWAR in this execution. We prove that $read_0(X_j)$ must perform a RAW or an AWAR in the execution fragment α_j. Suppose by contradiction that α_j does not contain a RAW or an AWAR.

The following claim shows that we can schedule a committed transaction T_j that writes a new value to X_j concurrent to $read_0(X_j)$ such that the execution is indistinguishable to both T_0 and T_j from a step contention-free execution (depicted in Figure 6.3a).

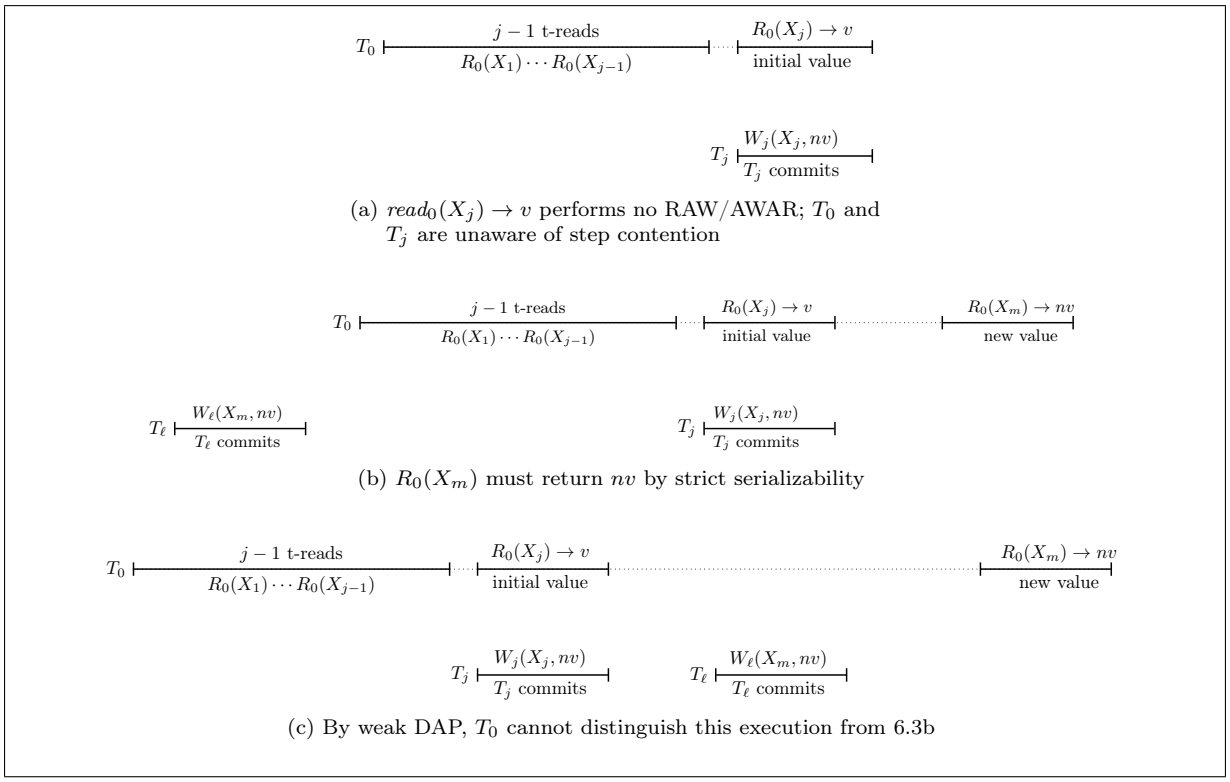

(a) $read_0(X_j) \to v$ performs no RAW/AWAR; T_0 and T_j are unaware of step contention

(b) $R_0(X_m)$ must return nv by strict serializability

(c) By weak DAP, T_0 cannot distinguish this execution from 6.3b

Figure 6.3: Executions in the proof of Theorem 6.3; execution in 6.3c is not strictly serializable

Claim 6.4. *For all $j \in \{1, \ldots, m-1\}$, M has an execution of the form $\alpha_1 \cdots \alpha_{j-1} \cdot \alpha_j^1 \cdot \delta_j \cdot \alpha_j^2$ where,*

- *δ_j is the t-complete step contention-free execution fragment of transaction T_j that writes $nv \neq v$ and commits*
- *$\alpha_j^1 \cdot \alpha_j^2 = \alpha_j$ is the complete execution fragment of the j^{th} t-read $read_0(X_j) \to v$ such that*
  - *α_j^1 does not contain any nontrivial events*
  - *$\alpha_1 \cdots \alpha_{j-1} \cdot \alpha_j^1 \cdot \delta_j \cdot \alpha_j^2$ is indistinguishable to T_0 from the step contention-free execution fragment $\alpha_1 \cdots \alpha_{j-1} \cdot \alpha_j^1 \cdot \alpha_j^2$*

Moreover, T_j does not access any base object to which T_0 applies a nontrivial event in $\alpha_1 \cdots \alpha_{j-1} \cdot \alpha_j^1 \cdot \delta_j$.

Proof. By wait-free progress (for read-only transactions) and strict serializability, M has an execution of the form $\alpha_1 \cdots \alpha_{j-1}$ in which each of the t-reads performed by T_0 must return the initial value of the t-objects.

Since T_j is an updating transaction, by sequential progress, there exists an execution of M of the form $\delta_j \cdot \alpha_1 \cdots \alpha_{j-1}$. Since T_0 and T_j are disjoint-access in the $\delta_j \cdot \alpha_1 \cdots \alpha_{j-1}$, by Lemma 2.10, T_0 and T_j do not contend on any base object in $\delta_j \cdot \alpha_1 \cdots \alpha_{j-1}$. Thus, $\alpha_1 \cdots \alpha_{j-1} \cdot \delta_j$ is indistinguishable to T_j from the execution δ_j and $\alpha_1 \cdots \alpha_{j-1} \cdot \delta_j$ is also an execution of M.

Let e be the first event that contains a write to a base object in α_j. If there exists no such write event to a base object in α_j, then $\alpha_j^1 = \alpha_j$ and α_j^2 is empty. Otherwise, we represent the execution fragment α_j as $\alpha_j^1 \cdot e \cdot \alpha_j^f$.

Since α_j^s does not contain any nontrivial events that write to a base object, $\alpha_1 \cdots \alpha_{j-1} \cdot \alpha_j^s \cdot \delta_j$ is indistinguishable to transaction T_j from the execution $\alpha_1 \cdots \alpha_{j-1} \cdot \delta_j$. Thus, $\alpha_1 \cdots \alpha_{j-1} \cdot \alpha_j^s \cdot \delta_j$ is an execution of M. Since e is not an atomic-write-after-read, $\alpha_1 \cdots \alpha_{j-1} \cdot \alpha_j^s \cdot \delta_j \cdot e$ is an execution of M. Since α_j does not contain a RAW, any read performed in α_j^f may only be performed to base objects

previously written in $e \cdot \alpha_j^f$. Thus, $\alpha_1 \cdots \alpha_{j-1} \cdot \alpha_j^s \cdot \delta_j \cdot e \cdot \alpha_j^f$ is indistinguishable to transaction T_0 from the step contention-free execution $\alpha_1 \cdots \alpha_{j-1} \cdot \alpha_j^s \cdot e \cdot \alpha_j^f$ in which $read_0(X_j) \to v$.

Choosing $\alpha_j^2 = e \cdot \alpha_j^f$, it follows that M has an execution of the form $\alpha_1 \cdots \alpha_{j-1} \cdot \alpha_j^1 \cdot \delta_j \cdot \alpha_j^2$ that is indistinguishable to T_j and T_0 from a step contention-free execution. The proof follows. □

We now prove that, for all $j \in \{1, \ldots, m-1\}$, M has an execution of the form $\delta_m \cdot \alpha_1 \cdots \alpha_{j-1} \cdot \alpha_j^1 \cdot \delta_j \cdot \alpha_j^2$ such that

- δ_m is the t-complete step contention-free execution of transaction T_ℓ that writes $nv \neq v$ to X_m and commits
- T_ℓ and T_0 do not contend on any base object in $\delta_m \cdot \alpha_1 \cdots \alpha_{j-1} \cdot \alpha_j^1 \cdot \delta_j \cdot \alpha_j^2$
- T_ℓ and T_j do not contend on any base object in $\delta_m \cdot \alpha_1 \cdots \alpha_{j-1} \cdot \alpha_j^1 \cdot \delta_j \cdot \alpha_j^2$.

By sequential progress for updating transactions, T_ℓ which writes the value nv to X_m must be committed in δ_m since it is running in the absence of step-contention from the initial configuration. Observe that T_ℓ and T_0 are disjoint-access in $\delta_m \cdot \alpha_1 \cdots \alpha_{j-1} \cdot \alpha_j^1 \cdot \delta_j \cdot \alpha_j^2$. By definition of α_j^1 and α_j^2, $\delta_m \cdot \alpha_1 \cdots \alpha_{j-1} \cdot \alpha_j^1 \cdot \delta_j \cdot \alpha_j^2$ is indistinguishable to T_0 from $\delta_m \cdot \alpha_1 \cdots \alpha_{j-1} \cdot \alpha_j^1 \cdot \alpha_j^2$. By Lemma 2.10, T_ℓ and T_0 do not contend on any base object in $\delta_m \cdot \alpha_1 \cdots \alpha_{j-1} \cdot \alpha_j^1 \cdot \alpha_j^2$.

By Claim 6.4, $\delta_m \cdot \alpha_1 \cdots \alpha_{j-1} \cdot \alpha_j^1 \cdot \delta_j$ is indistinguishable to T_j from $\delta_m \cdot \delta_j$. But transactions T_ℓ and T_j are disjoint-access in $\delta_m \cdot \delta_j$, and by Lemma 2.10, T_j and T_ℓ do not contend on any base object in $\delta_m \cdot \delta_j$.

Since strict serializability of M stipulates that each of the j t-reads performed by T_0 return the initial values of the respective t-objects, M has an execution of the form $\delta_m \cdot \alpha_1 \cdots \alpha_{j-1} \cdot \alpha_j^1 \cdot \delta_j \cdot \alpha_j^2$.

Consider the extension of $\delta_m \cdot \alpha_1 \cdots \alpha_{j-1} \cdot \alpha_j^1 \cdot \delta_j \cdot \alpha_j^2$ in which T_0 performs $(m-j)$ t-reads of $X_{j+1}, \cdots, X_m$ step contention-free and commits (depicted in Figure 6.3b). By wait-free progress of M and since T_0 is a read-only transaction, there exists such an execution. Notice that the m^{th} t-read, $read_0(X_m)$ must return the value nv by strict serializability since T_ℓ precedes T_0 in real-time order in this execution.

Recall that neither pairs of transactions T_ℓ and T_j nor T_ℓ and T_0 contend on any base object in the execution $\delta_m \cdot \alpha_1 \cdots \alpha_{j-1} \cdot \alpha_j^1 \cdot \delta_j \cdot \alpha_j^2$. It follows that for all $j \in \{1, \ldots, m-1\}$, M has an execution of the form $\alpha_1 \cdots \alpha_{j-1} \cdot \alpha_j^1 \cdot \delta_j \cdot \alpha_j^2 \cdot \delta_m$ in which T_j precedes T_ℓ in real-time order.

Let α' be the execution fragment that extends $\alpha_1 \cdots \alpha_{j-1} \cdot \alpha_j^1 \cdot \delta_j \cdot \alpha_j^2 \cdot \delta_m$ in which T_0 performs $(m-j)$ t-reads of $X_{j+1}, \cdots, X_m$ step contention-free and commits (depicted in Figure 6.3c). Since $\alpha_1 \cdots \alpha_{j-1} \cdot \alpha_j^1 \cdot \delta_j \cdot \alpha_j^2 \cdot \delta_m$ is indistinguishable to T_0 from the execution $\delta_m \cdot \alpha_1 \cdots \alpha!_{j-1} \cdot \alpha_j^1 \cdot \delta_j \cdot \alpha_j^2$, $read_0(X_m)$ must return the response value nv in α'.

The execution $\alpha_1 \cdots \alpha_{j-1} \cdot \alpha_j^1 \cdot \delta_j \cdot \alpha_j^2 \cdot \delta_m \cdot \alpha'$ is not strictly serializable. In any serialization, T_j must precede T_ℓ to respect the real-time ordering of transactions, while T_ℓ must precede T_0 since $read_j(X_m)$ returns the value of X_m updated by T_ℓ. Also, transaction T_0 must precede T_j since $read_0(X_j)$ returns the initial value of X_j. But there exists no such serialization: a contradiction to the assumption that M is strict serializable.

Thus, for all $j \in \{1, \ldots, m-1\}$, transaction T_0 must perform a RAW or an AWAR during the execution of $read_0(X_j)$, completing the proof. □ □

Since Theorem 6.3 implies that read-only transactions must perform nontrivial events, we have the following corollary that was proved directly in [24].

Corollary 6.5 ([24])**.** *There does not exist any strictly serializable weak DAP TM implementation $M \in \mathcal{RWF}$ that uses invisible reads.*

6.5 Related work and Discussion

Attiya *et al.* [24] showed that it is impossible to implement weak DAP strictly serializable TMs in $\mathcal{RWF}$ if read-only transactions may only apply trivial primitives to base objects. Attiya et al. [24] also considered a stronger "disjoint-access" property, called simply DAP, referring to the original definition proposed Israeli and Rappoport [86]. In DAP, two transactions are allowed to *concurrently access* (even for reading) the same base object only if they are disjoint-access. For an n-process DAP TM implementation, it is shown in [24] that a read-only transaction must perform at least $n - 3$ writes. Our lower bound in Theorem 6.3 is strictly stronger than the one in [24], as it assumes only weak DAP, considers a more precise RAW/AWAR metric, and does not depend on the number of processes in the system. (Technically, the last point follows from the fact that the execution constructed in the proof of Theorem 6.3 uses only 3 concurrent processes.) Thus, the theorem subsumes the two lower bounds of [24] within a single proof.

Perelman *et al.* [110] considered the closely related (to $\mathcal{RWF}$) class of *mv-permissive* TMs: a transaction can only be aborted if it is an updating transaction that conflicts with another updating transaction. $\mathcal{RWF}$ is incomparable with the class of mv-permissive TMs. On the one hand, mv-permissiveness guarantees that read-only transactions never abort, but does not imply that they commit in a wait-free manner. On the other hand, $\mathcal{RWF}$ allows an updating transaction to abort in the presence of a concurrent read-only transaction, which is disallowed by mv-permissive TMs. Observe that, technically, mv-permissiveness is a blocking TM-progress condition, although when used in conjunction with wait-free TM-liveness, it is a partially non-blocking TM-progress condition that is strictly stronger than $\mathcal{RWF}$.

Assuming starvation-free TM-liveness, [110] showed that implementing a weak DAP strictly serializable mv-permissive TM is impossible. In the thesis, we showed that strictly serializable TMs in $\mathcal{RWF}$ cannot provide strict DAP, but proving the impossibility result assuming weak DAP remains an interesting open question.

[110] also proved that mv-permissive TMs cannot be *online space optimal*, *i.e.*, no mv-permissive TM can keep the minimum number of old object versions for any TM history. Our result on the space complexity of implementations in $\mathcal{RWF}$ that use invisible reads (Theorem 6.1) is different since it proves that the implementation must maintain an unbounded number of versions of every t-object. Our proof technique can however be used to show that mv-permissive TMs considered in [110] should also maintain unbounded number of versions.

7

Hybrid transactional memory (HyTM)

HAL: The 9000 series is the most reliable computer ever made. No 9000 computer has ever made a mistake or distorted information. We are all, by any practical definition of the words, foolproof and incapable of error.
. . .
HAL: I've just picked up a fault in the AE35 unit. It's going to go 100% failure in 72 hours.
HAL: It can only be attributable to human error.

Stanley Kubrick-2001: A Space Odyssey

7.1 Overview

Hybrid transactional memory. The TM abstraction, in its original manifestation from the proposal by Herlihy and Moss [80], augmented the processor's *cache-coherence protocol* and extended the CPU's instruction set with instructions to indicate which memory accesses must be transactional [80]. Most popular TM designs, subsequent to the original proposal in [80] have implemented all the functionality in software [36, 52, 79, 101, 117] (cf. software TM model in Chapter 2). More recently, CPUs have included hardware extensions to support *short*, *small* hardware transactions [1, 107, 111].

Early experience with programming *Hardware transactional memory (HTM)*, *e.g.* [7, 38, 44], paints an interesting picture: if used carefully, HTM can be an extremely useful construct, and can significantly speed up and simplify concurrent implementations. At the same time, this powerful tool is not without its limitations: since HTMs are usually implemented on top of the cache coherence mechanism, hardware transactions have inherent *capacity constraints* on the number of distinct memory locations that can be accessed inside a single transaction. Moreover, all current proposals are *best-effort*, as they may abort

under imprecisely specified conditions (cache capacity overflow, interrupts *etc*). In brief, the programmer should not solely rely on HTMs.

Several *Hybrid Transactional Memory (HyTM)* schemes [35, 37, 88, 99] have been proposed to complement the fast, but best-effort nature of HTM with a slow, reliable software transactional memory (STM) backup. These proposals have explored a wide range of trade-offs between the overhead on hardware transactions, concurrent execution of hardware and software, and the provided progress guarantees.

Early proposals for HyTM implementations [37, 88] shared some interesting features. First, transactions that do not conflict are expected to run concurrently, regardless of their types (software or hardware). This property is referred to as *progressiveness* [63] and is believed to allow for increased parallelism. Second, in addition to changing the values of transactional objects, hardware transactions usually employ *code instrumentation* techniques. Intuitively, instrumentation is used by hardware transactions to detect concurrency scenarios and abort in the case of contention. The number of instrumentation steps performed by these implementations within a hardware transaction is usually proportional to the size of the transaction's data set.

Recent work by Riegel *et al.* [113] surveyed the various HyTM algorithms to date, focusing on techniques to reduce instrumentation overheads in the frequently executed hardware fast-path. However, it is not clear whether there are fundamental limitations when building a HyTM with non-trivial concurrency between hardware and software transactions. In particular, what are the inherent instrumentation costs of building a HyTM, and what are the trade-offs between these costs and the provided *concurrency*, *i.e.*, the ability of the HyTM system to run software and hardware transactions in parallel?

Modelling HyTM. To address these questions, the thesis proposes the first model for hybrid TM systems which formally captures the notion of *cached* accesses provided by hardware transactions, and precisely defines instrumentation costs in a quantifiable way.

We model a hardware transaction as a series of memory accesses that operate on locally cached copies of the variables, followed by a *cache-commit* operation. In case a concurrent transaction performs a (read-write or write-write) conflicting access to a cached object, the cached copy is invalidated and the hardware transaction aborts.

Our model for instrumentation is motivated by recent experimental evidence which suggests that the overhead on hardware transactions imposed by code which detects concurrent software transactions is a significant performance bottleneck [102]. In particular, we say that a HyTM implementation imposes a logical partitioning of shared memory into *data* and *metadata* locations. Intuitively, metadata is used by transactions to exchange information about contention and conflicts while data locations only store the *values* of data items read and updated within transactions. We quantify instrumentation cost by measuring the number of accesses to *metadata objects* which transactions perform. Our framework captures all known HyTM proposals which combine HTMs with an STM fallback [35, 37, 88, 99, 112].

The cost of instrumentation. Once this general model is in place, we derive two lower bounds on the cost of implementing a HyTM. First, we show that some instrumentation is necessary in a HyTM implementation even if we only intend to provide *sequential* progress, where a transaction is only guaranteed to commit if it runs in the absence of concurrency.

Second, we prove that any progressive HyTM implementation providing *obstruction-free liveness* (every operation running *solo* returns some response) and has executions in which an arbitrarily long read-only hardware transaction running in the absence of concurrency *must* access a number of distinct metadata objects proportional to the size of its data set. Our proof technique is interesting in its own right. Inductively, we start with a sequential execution in which a "large" set S_m of read-only hardware transactions, each accessing m distinct data items and m distinct metadata memory locations, run after an execution E_m. We then construct execution E_{m+1}, an extension of E_m, which forces at least half of the transactions in S_m to access a *new* metadata base object when reading a new $(m+1)^{th}$ data item, running after E_{m+1}. The technical challenge, and the key departure from prior work on STM lower bounds, *e.g.* [24, 60, 64], is that hardware transactions practically possess "automatic" conflict detection, aborting on contention. This is in contrast to STMs, which must take steps to detect contention on memory locations.

We match this lower bound with an HyTM algorithm that, additionally, allows for uninstrumented writes and *invisible reads* and is provably *opaque* [64]. To the best of our knowledge, this is the first formal proof of correctness of a HyTM algorithm.

Low-instrumentation HyTM. The high instrumentation costs of early HyTM designs, which we show to be inherent, stimulated more recent HyTM schemes [35, 99, 102, 113] to sacrifice progressiveness for *constant* instrumentation cost (*i.e.*, not depending on the size of the transaction). In the past two years, Dalessandro *et al.* [35] and Riegel *et al.* [113] have proposed HyTMs based on the efficient *NOrec STM* [36]. These HyTMs schemes do not guarantee any parallelism among transactions; only sequential progress is ensured. Despite this, they are among the best-performing HyTMs to date due to the limited instrumentation in hardware transactions.

Starting from this observation, we provide a more precise upper bound for *low-instrumentation* HyTMs by presenting a HyTM algorithm with invisible reads *and* uninstrumented hardware writes which guarantees that a hardware transaction accesses at most one metadata object in the course of its execution. Software transactions in this implementation remain progressive, while hardware transactions are guaranteed to commit only if they do not run concurrently with an updating software transaction (or exceed capacity). Therefore, the cost of avoiding the linear lower bound for progressive implementations is that hardware transactions may be aborted by non-conflicting software ones.

Roadmap of Chapter 7. In Section 7.2, we introduce the model of HyTMs and Section 7.3 studies the inherent cost of concurrency in progressive HyTMs by presenting a linear lower bound on the cost of instrumentation while Section 7.4 presents a matching upper bound. In Section 7.5 discusses providing partial concurrency with instrumentation cost and in Section 7.6, we elaborate on prior work related to HyTMs.

7.2 Modelling HyTM

In this chapter, we introduce the model of HyTMs, extending the TM model from Chapter 2, that intuitively captures the cache-coherence protocols employed in shared memory systems.

7.2.1 Direct and cached accesses

We now describe the operation of a *Hybrid Transactional Memory (HyTM)* implementation. In our model, every base object can be accessed with two kinds of primitives, *direct* and *cached*.

In a direct access, the rmw primitive operates on the memory state: the direct-access event atomically reads the value of the object in the shared memory and, if necessary, modifies it.

In a cached access performed by a process i, the rmw primitive operates on the *cached* state recorded in process i's *tracking set* τ_i. One can think of τ_i as the *L1 cache* of process i. A *hardware transaction* is a series of cached rmw primitives performed on τ_i followed by a *cache-commit* primitive.

More precisely, τ_i is a set of triples (b, v, m) where b is a base object identifier, v is a value, and $m \in \{shared, exclusive\}$ is an access *mode*. The triple (b, v, m) is added to the tracking set when i performs a cached rmw access of b, where m is set to *exclusive* if the access is nontrivial, and to *shared* otherwise. We assume that there exists some constant TS (representing the size of the L1 cache) such that the condition $|\tau_i| \leq TS$ must always hold; this condition will be enforced by our model. A base object b is *present* in τ_i with mode m if $\exists v, (b, v, m) \in \tau_i$.

A trivial (resp. nontrivial) cached primitive $\langle g, h\rangle$ applied to b by process i first checks the condition $|\tau_i| = TS$ and if so, it sets $\tau_i = \emptyset$ and immediately returns $\perp$ (we call this event a *capacity abort*). We assume that TS is large enough so that no transaction with data set of size 1 can incur a capacity abort. If the transaction does not incur a capacity abort, the process checks whether b is present in exclusive (resp. any) mode in τ_j for any $j \neq i$. If so, τ_i is set to $\emptyset$ and the primitive returns $\perp$. Otherwise, the triple $(b, v, shared)$ (resp. $(b, g(v), exclusive)$) is added to τ_i, where v is the most recent cached value of

(a) τ_2 is invalidated by (fast-path or slow-path) transaction T_1's access of base object b

(b) τ_2 is invalidated by (fast-path or slow-path) transaction T_1's write to base object b

Figure 7.1: Tracking set aborts in fast-path transactions; we denote a fast-path (and resp. slow-path) transaction by F (and resp. S)

b in τ_i (in case b was previously accessed by i within the current hardware transaction) or the value of b in the current memory configuration, and finally $h(v)$ is returned.

A tracking set can be *invalidated* by a concurrent process: if, in a configuration C where $(b, v, exclusive) \in \tau_i$ (resp. $(b, v, shared) \in \tau_i$), a process $j \neq i$ applies any primitive (resp. any *nontrivial* primitive) to b, then τ_i becomes *invalid* and any subsequent cached primitive invoked by i sets τ_i to $\emptyset$ and returns $\perp$. We refer to this event as a *tracking set abort*.

Finally, the *cache-commit* primitive issued by process i with a valid τ_i does the following: for each base object b such that $(b, v, exclusive) \in \tau_i$, the value of b in C is updated to v. Finally, τ_i is set to $\emptyset$ and the primitive returns *commit*.

Note that HTM may also abort spuriously, or because of unsupported operations [111]. The first cause can be modelled probabilistically in the above framework, which would not however significantly affect our claims and proofs, except for a more cumbersome presentation. Also, our lower bounds are based exclusively on executions containing t-reads and t-writes. Therefore, in the following, we only consider contention and capacity aborts.

7.2.2 Slow-path and fast-path transactions

In the following, we partition HyTM transactions into *fast-path transactions* and *slow-path transactions*. Practically, two separate algorithms (fast-path one and slow-path one) are provided for each t-operation.

A slow-path transaction models a regular software transaction. An event of a slow-path transaction is either an invocation or response of a t-operation, or a rmw primitive on a base object.

A fast-path transaction essentially encapsulates a hardware transaction. An event of a fast-path transaction is either an invocation or response of a t-operation, a cached primitive on a base object, or a *cache-commit*: *t-read* and *t-write* are only allowed to contain cached primitives, and *tryC* consists of invoking *cache-commit*. Furthermore, we assume that a fast-path transaction T_k returns A_k as soon an underlying cached primitive or *cache-commit* returns $\perp$. Figure 7.1 depicts such a scenario illustrating a tracking set abort: fast-path transaction T_2 executed by process p_2 accesses a base object b in shared (and resp. exclusive) mode and it is added to its tracking set τ_2. Immediately after the access of b by T_2, a concurrent transaction T_1 applies a nontrivial primitive to b (and resp. accesses b). Thus, the tracking of p_2 is invalidated and T_2 must be aborted in any extension of this execution.

We provide two key observations on this model regarding the interactions of non-committed fast path transactions with other transactions. Let E be any execution of a HyTM implementation $\mathcal{M}$ in which a fast-path transaction T_k is either t-incomplete or aborted. Then the sequence of events E' derived by removing all events of $E|k$ from E is an execution $\mathcal{M}$. Moreover:

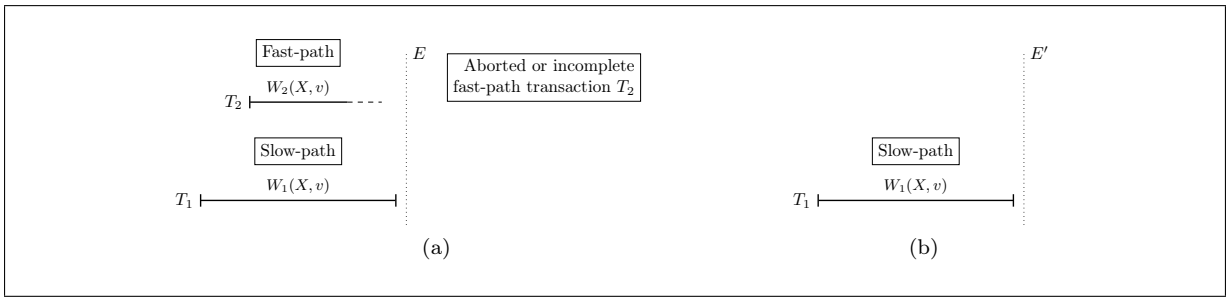

Figure 7.2: Execution E in Figure 7.2a is indistinguishable to T_1 from the execution E' in Figure 7.2b

Observation 7.1. *To every slow-path transaction $T_m \in txns(E)$, E is indistinguishable from E'.*

Observation 7.2. *If a fast-path transaction $T_m \in txns(E) \setminus \{T_k\}$ does not incur a tracking set abort in E, then E is indistinguishable to T_m from E'.*

Intuitively, these observations say that fast-path transactions which are not yet committed are invisible to slow-path transactions, and can communicate with other fast-path transactions only by incurring their tracking-set aborts. Figure 7.2 illustrates Observation 7.1: a fast-path transaction T_2 is concurrent to a slow-path transaction T_1 in an execution E. Since T_2 is t-incomplete or aborted in this execution, E is indistinguishable to T_1 from an execution E' derived by removing all events of T_2 from E. Analogously, to illustrate Observation 7.2, if T_1 is a fast-path transaction that does not incur a tracking set abort in E, then E is indistinguishable to T_1 from E'.

7.2.3 Instrumentation

Now we define the notion of *code instrumentation* in fast-path transactions. Intuitively, instrumentation characterizes the number of extra "metadata" accesses performed by a fast-path transaction.

We start with the following technical definition. An execution E of a HyTM $\mathcal{M}$ *appears t-sequential* to a transaction $T_k \in txns(E)$ if there exists an execution E' of $\mathcal{M}$ such that:

- $txns(E') \subseteq txns(E) \setminus \{T_k\}$ and the configuration after E' is t-quiescent,
- every transaction $T_m \in txns(E)$ that precedes T_k in real-time order is included in E' such that $E|m = E'|m$,
- for every transaction $T_m \in txns(E')$, $Rset_{E'}(T_m) \subseteq Rset_E(T_m)$ and $Wset_{E'}(T_m) \subseteq Wset_E(T_m)$, and
- $E' \cdot E|k$ is an execution of $\mathcal{M}$.

Definition 7.1 (Data and metadata base objects)**.** *Let $\mathcal{X}$ be the set of t-objects operated by a HyTM implementation $\mathcal{M}$. Now we partition the set of base objects used by $\mathcal{M}$ into a set $\mathbb{D}$ of* data *objects and a set $\mathbb{M}$ of* metadata *objects ($\mathbb{D} \cap \mathbb{M} = \emptyset$). We further partition $\mathbb{D}$ into sets $\mathbb{D}_X$ associated with each t-object $X \in \mathcal{X}$: $\mathbb{D} = \bigcup_{X \in \mathcal{X}} \mathbb{D}_X$, for all $X \neq Y$ in $\mathcal{X}$, $\mathbb{D}_X \cap \mathbb{D}_Y = \emptyset$, such that:*

1. *In every execution E, each fast-path transaction $T_k \in txns(E)$ only accesses base objects in $\bigcup_{X \in DSet(T_k)} \mathbb{D}_X$ or $\mathbb{M}$.*
2. *Let $E \cdot \rho$ and $E \cdot E' \cdot \rho'$ be two t-complete executions, such that E and $E \cdot E'$ are t-complete, ρ and ρ' are complete executions of a transaction $T_k \notin txns(E \cdot E')$, $H_\rho = H_{\rho'}$, and $\forall T_m \in txns(E')$, $Dset(T_m) \cap Dset(T_k) = \emptyset$. Then the states of the base objects $\bigcup_{X \in DSet(T_k)} \mathbb{D}_X$ in the configuration after $E \cdot \rho$ and $E \cdot E' \cdot \rho'$ are the same.*

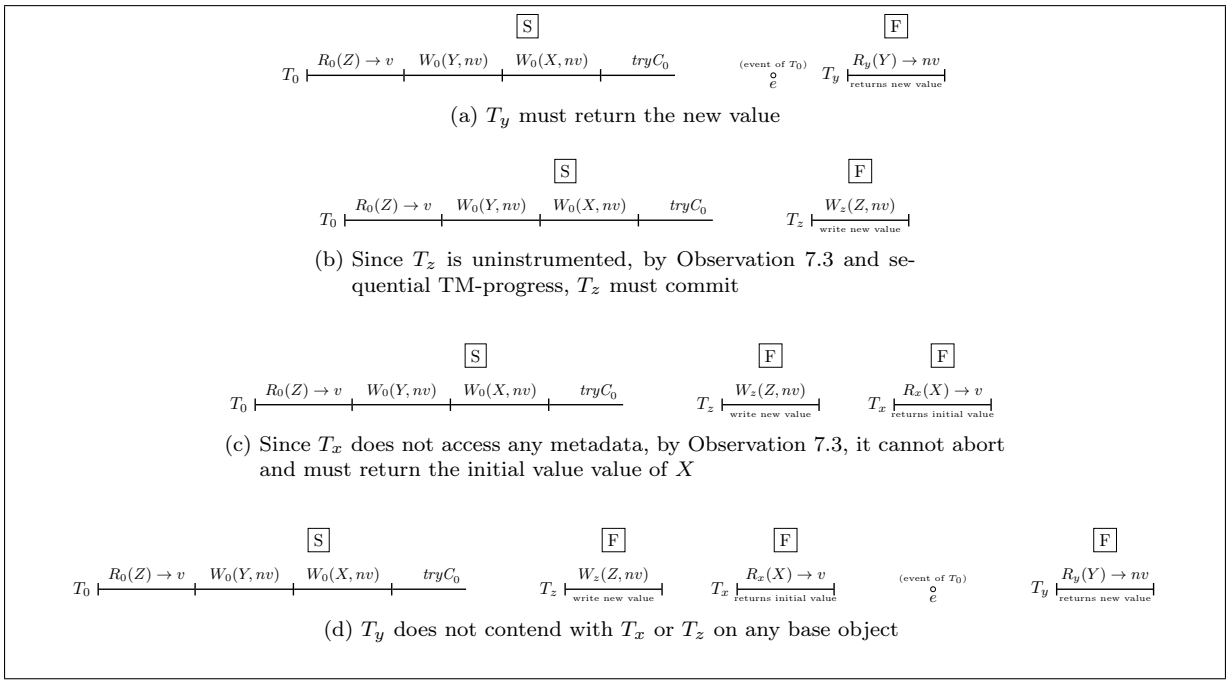
S
F
T_0 $R_0(Z) \to v$ $W_0(Y, nv)$ $W_0(X, nv)$ $tryC_0$
(event of T_0) e
T_y $R_y(Y) \to nv$ returns new value
S
F
T_0 $R_0(Z) \to v$ $W_0(Y, nv)$ $W_0(X, nv)$ $tryC_0$
T_z $W_z(Z, nv)$ write new value
S
F
F
T_0 $R_0(Z) \to v$ $W_0(Y, nv)$ $W_0(X, nv)$ $tryC_0$
T_z $W_z(Z, nv)$ write new value
T_x $R_x(X) \to v$ returns initial value
S
F
F
F
T_0 $R_0(Z) \to v$ $W_0(Y, nv)$ $W_0(X, nv)$ $tryC_0$
T_z $W_z(Z, nv)$ write new value
T_x $R_x(X) \to v$ returns initial value
(event of T_0) e
T_y $R_y(Y) \to nv$ returns new value

(a) T_y must return the new value

(b) Since T_z is uninstrumented, by Observation 7.3 and sequential TM-progress, T_z must commit

(c) Since T_x does not access any metadata, by Observation 7.3, it cannot abort and must return the initial value value of X

(d) T_y does not contend with T_x or T_z on any base object

Figure 7.3: Executions in the proof of Theorem 7.4; execution in 7.3d is not strictly serializable

3. *Let execution E appear t-sequential to a transaction T_k and let the enabled event e of T_k after E be a primitive on a base object $b \in \mathbb{D}$. Then, unless e returns $\bot$, $E \cdot e$ also appears t-sequential to T_k.*

Intuitively, the first condition says that a transaction is only allowed to access data objects based on its data set. The second condition says that transactions with disjoint data sets can communicate only via metadata objects. Finally, the last condition means that base objects in $\mathbb{D}$ may only contain the "values" of t-objects, and cannot be used to detect concurrent transactions. Note that our results will lower bound the number of metadata objects that must be accessed under particular assumptions, thus from a cost perspective, $\mathbb{D}$ should be made as large as possible.

All HyTM proposals we aware of, such as *HybridNOrec* [35, 112], *PhTM* [99] and others [37, 88], conform to our definition of instrumentation in fast-path transactions. For instance, HybridNOrec [35, 112] employs a distinct base object in $\mathbb{D}$ for each t-object and a global *sequence lock* as the metadata that is accessed by fast-path transactions to detect concurrency with slow-path transactions. Similarly, the HyTM implementation by *Damron et al.* [37] also associates a distinct base object in $\mathbb{D}$ for each t-object and additionally, a *transaction header* and *ownership record* as metadata base objects.

Definition 7.2 (Uninstrumented HyTMs)**.** *A HyTM implementation $\mathcal{M}$ provides* uninstrumented writes (resp. reads) *if in every execution E of $\mathcal{M}$, for every write-only (resp. read-only) fast-path transaction T_k, all primitives in $E|k$ are performed on base objects in $\mathbb{D}$. A HyTM is uninstrumented if both its reads and writes are uninstrumented.*

Observation 7.3. *Consider any execution E of a HyTM implementation $\mathcal{M}$ which provides uninstrumented reads (resp. writes). For any fast-path read-only (resp. write-only) transaction $T_k \notin txns(E)$, that runs step-contention free after E, the execution E appears t-sequential to T_k.*

7.2.4 Impossibility of uninstrumented HyTMs

In this section, we show that any strictly serializable HyTM must be instrumented, even under a very weak progress assumption by which a transaction is guaranteed to commit only when run t-sequentially:

Definition 7.3 (Sequential TM-progress for HyTMs)**.** *A HyTM implementation $\mathcal{M}$ provides* sequential TM-progress for fast-path transactions (and resp. slow-path) *if in every execution E of $\mathcal{M}$, a fast-path (and resp. slow-path) transaction T_k returns A_k in E only if T_k incurs a capacity abort or T_k is concurrent to another transaction. We say that $\mathcal{M}$ provides sequential TM-progress if it provides sequential TM-progress for fast-path and slow-path transactions.*

Theorem 7.4. *There does not exist a strictly serializable uninstrumented HyTM implementation that ensures sequential TM-progress and TM-liveness.*

Proof. Suppose by contradiction that such a HyTM $\mathcal{M}$ exists. For simplicity, assume that v is the initial value of t-objects X, Y and Z. Let E be the t-complete step contention-free execution of a slow-path transaction T_0 that performs $read_0(Z) \to v$, $write_0(X, nv)$, $write_0(Y, nv)$ ($nv \neq v$), and commits. Such an execution exists since $\mathcal{M}$ ensures sequential TM-progress.

By Observation 7.3, any transaction that runs step contention-free starting from a prefix of E must return a non-abort value. Since any such transaction reading X or Y must return v when it starts from the empty prefix of E and nv when it starts from E.

Thus, there exists E', the longest prefix of E that cannot be extended with the t-complete step contention-free execution of a *fast-path* transaction reading X or Y and returning nv. Let e is the enabled event of T_0 in the configuration after E'. Without loss of generality, suppose that there exists an execution $E' \cdot e \cdot E_y$ where E_y is the t-complete step contention-free execution fragment of some fast-path transaction T_y that reads Y is returns nv (Figure 7.3a).

Claim 7.5. *$\mathcal{M}$ has an execution $E' \cdot E_z \cdot E_x$, where*

- *E_z is the t-complete step contention-free execution fragment of a fast-path transaction T_z that writes $nv \neq v$ to Z and commits*
- *E_x is the t-complete step contention-free execution fragment of a fast-path transaction T_x that performs a single t-read $read_x(X) \to v$ and commits.*

Proof. By Observation 7.3, the extension of E' in which T_z writes to Z and tries to commit appears t-sequential to T_z. By sequential TM-progress, T_z complets the write and commits. Let $E' \cdot E_z$ (Figure 7.3b) be the resulting execution of $\mathcal{M}$.

Similarly, the extension of E' in which T_x reads X and tries to commit appears t-sequential to T_x. By sequential TM-progress, T_x commits and let $E' \cdot E_x$ be the resulting execution of $\mathcal{M}$. By the definition of E', $read_x(X)$ must return v in $E' \cdot E_x$.

Since $\mathcal{M}$ is uninstrumented and the data sets of T_x and T_z are disjoint, the sets of base objects accessed in the execution fragments E_x and E_y are also disjoint. Thus, $E' \cdot E_z \cdot E_x$ is indistinguishable to T_x from the execution $E' \cdot E_x$, which implies that $E' \cdot E_z \cdot E_x$ is an execution of $\mathcal{M}$ (Figure 7.3c). □

Finally, we prove that the sequence of events, $E' \cdot E_z \cdot E_x \cdot e \cdot E_y$ is an execution of $\mathcal{M}$.

Since the transactions T_x, T_y, T_z have pairwise disjoint data sets in $E' \cdot E_z \cdot E_x \cdot e \cdot E_y$, no base object accessed ib E_y can be accessed in E_x and E_z. The read operation on X performed by T_y in $E' \cdot e \cdot E_y$ returns nv and, by the definition of E' and e, T_y must have accessed the base object b modified in the event e by T_0. Thus, b is not accessed in E_x and E_z and $E' \cdot E_z \cdot E_x \cdot e$ is an execution of $\mathcal{M}$. Summing up, $E' \cdot E_z \cdot E_x \cdot e \cdot E_y$ is indistinguishable to T_y from $E' \cdot e \cdot E_y$, which implies that $E' \cdot E_z \cdot E_x \cdot e \cdot E_y$ is an execution of $\mathcal{M}$ (Figure 7.3d).

But the resulting execution is not strictly serializable. Indeed, suppose that a serialization exists. As the value written by T_0 is returned by a committed transaction T_y, T_0 must be committed and precede T_y in the serialization. Since T_x returns the initial value of X, T_x must precede T_0. Since T_0 reads the initial value of Z, T_0 must precede T_z. Finally, T_z must precede T_x to respect the real-time order. The cycle in the serialization establishes a contradiction. □

7.3 A linear lower bound on instrumentation for progressive HyTMs

In this section, we show that giving HyTM the ability to run and commit transactions in parallel brings considerable instrumentation costs. We focus on a natural progress condition called progressiveness [61, 62, 63] that allows a transaction to abort only if it experiences a read-write or write-write conflict with a concurrent transaction:

Definition 7.4 (Progressiveness for HyTMs)**.** *We say that transactions T_i and T_j* conflict *in an execution E on a t-object X if $X \in Dset(T_i) \cap Dset(T_j)$ and $X \in Wset(T_i) \cup Wset(T_j)$.*

A HyTM implementation $\mathcal{M}$ is fast-path *(resp.* slow-path*)* progressive *if in every execution E of $\mathcal{M}$ and for every fast-path (and resp. slow-path) transaction T_i that aborts in E, either A_i is a capacity abort or T_i conflicts with some transaction T_j that is concurrent to T_i in E. We say $\mathcal{M}$ is* progressive *if it is both fast-path and slow-path progressive.*

We show that for every opaque fast-path progressive HyTM that provides obstruction-free TM-liveness, an arbitrarily long read-only transaction might access a number of distinct metadata base objects that is linear in the size of its read set or experience a capacity abort.

The following auxiliary results will be crucial in proving our lower bound. We observe first that a fast path transaction in a progressive HyTM can contend on a base object only with a conflicting transaction.

Lemma 7.6. *Let $\mathcal{M}$ be any fast-path progressive HyTM implementation. Let $E \cdot E_1 \cdot E_2$ be an execution of $\mathcal{M}$ where E_1 (and resp. E_2) is the step contention-free execution fragment of transaction $T_1 \notin txns(E)$ (and resp. $T_2 \notin txns(E)$), T_1 (and resp. T_2) does not conflict with any transaction in $E \cdot E_1 \cdot E_2$, and at least one of T_1 or T_2 is a fast-path transaction. Then, T_1 and T_2 do not contend on any base object in $E \cdot E_1 \cdot E_2$.*

Proof. Suppose, by contradiction that T_1 or T_2 contend on the same base object in $E \cdot E_1 \cdot E_2$.

If in E_1, T_1 performs a nontrivial event on a base object on which they contend, let e_1 be the last event in E_1 in which T_1 performs such an event to some base object b and e_2, the first event in E_2 that accesses b. Otherwise, T_1 only performs trivial events in E_1 to base objects on which it contends with T_2 in $E \cdot E_1 \cdot E_2$: let e_2 be the first event in E_2 in which E_2 performs a nontrivial event to some base object b on which they contend and e_1, the last event of E_1 in T_1 that accesses b.

Let E_1' (and resp. E_2') be the longest prefix of E_1 (and resp. E_2) that does not include e_1 (and resp. e_2). Since before accessing b, the execution is step contention-free for T_1, $E \cdot E_1' \cdot E_2'$ is an execution of $\mathcal{M}$. By construction, T_1 and T_2 do not conflict in $E \cdot E_1' \cdot E_2'$. Moreover, $E \cdot E_1 \cdot E_2'$ is indistinguishable to T_2 from $E \cdot E_1' \cdot E_2'$. Hence, T_1 and T_2 are poised to apply contending events e_1 and e_2 on b in the execution $\tilde{E} = E \cdot E_1' \cdot E_2'$. Recall that at least one event of e_1 and e_2 must be nontrivial.

Consider the execution $\tilde{E} \cdot e_1 \cdot e_2'$ where e_2' is the event of p_2 in which it applies the primitive of e_2 to the configuration after $\tilde{E} \cdot e_1$. After $\tilde{E} \cdot e_1$, b is contained in the tracking set of process p_1. If b is contained in τ_1 in the shared mode, then e_2' is a nontrivial primitive on b, which invalidates τ_1 in $\tilde{E} \cdot e_1 \cdot e_2'$. If b is contained in τ_1 in the exclusive mode, then any subsequent access of b invalidates τ_1 in $\tilde{E} \cdot e_1 \cdot e_2'$. In both cases, τ_1 is invalidated and T_1 incurs a tracking set abort. Thus, transaction T_1 must return A_1 in any extension of $E \cdot e_1 \cdot e_2$—a contradiction to the assumption that $\mathcal{M}$ is progressive. □

Iterative application of Lemma 7.6 implies the following:

Corollary 7.7. *Let $\mathcal{M}$ be any fast-path progressive HyTM implementation. Let $E \cdot E_1 \cdots E_i \cdot E_{i+1} \cdots E_m$ be any execution of $\mathcal{M}$ where for all $i \in \{1, \ldots, m\}$, E_i is the step contention-free execution fragment of transaction $T_i \notin txns(E)$ and any two transactions in $E_1 \cdots E_m$ do not conflict. For all $i, j = 1, \ldots, m$, $i \neq j$, if T_i is fast-path, then T_i and T_j do not contend on a base object in $E \cdot E_1 \cdots E_i \cdots E_m$*

Proof. Let T_i be a fast-path transaction. By Lemma 7.6, in $E \cdot E_1 \cdots E_i \cdots E_m$, T_i does not contend with T_{i-1} (if $i > 1$) or T_{i+1} (if $i < m$) on any base object and, thus, E_i commutes with E_{i-1} and E_{i+1}. Thus, $E \cdot E_1 \cdots E_{i-2} \cdot E_i \cdot E_{i-1} \cdot E_{i+1} \cdots E_m$ (if $i > 1$) and $E \cdot E_1 \cdots E_{i-1} \cdot E_{i+1} \cdot E_i \cdot E_{i+2} \cdots E_m$ (if $i < m$) are executions of $\mathcal{M}$. By iteratively applying Lemma 7.6, we derive that T_i does not contend with any T_j, $j \neq i$. □

We say that execution fragments E and E' are *similar* if they export equivalent histories, *i.e.*, no process can see the difference between them by looking at the invocations and responses of t-operations. We now use Corollary 7.7 to show that t-operations only accessing data base objects cannot detect contention with non-conflicting transactions.

Lemma 7.8. *Let E be any t-complete execution of a progressive HyTM implementation $\mathcal{M}$ that provides OF TM-liveness. For any $m \in \mathbb{N}$, consider a set of m executions of $\mathcal{M}$ of the form $E \cdot E_i \cdot \gamma_i \cdot \rho_i$ where E_i is the t-complete step contention-free execution fragment of a transaction T_{m+i}, γ_i is a complete step contention-free execution fragment of a* fast-path *transaction T_i such that $Dset(T_i) \cap Dset(T_{m+i}) = \emptyset$ in $E \cdot E_i \cdot \gamma_i$, and ρ_i is the execution fragment of a t-operation by T_i that does not contain accesses to any metadata base object. If, for all $i, j \in \{1, \ldots, m\}$, $i \neq j$, $Dset(T_i) \cap Dset(T_{m+j}) = \emptyset$, $Dset(T_i) \cap Dset(T_j) = \emptyset$ and $Dset(T_{m+i}) \cap Dset(T_{m+j}) = \emptyset$, then there exists a t-complete step contention-free execution fragment E' that is similar to $E_1 \cdots E_m$ such that for all $i \in \{1, \ldots, m\}$, $E \cdot E' \cdot \gamma_i \cdot \rho_i$ is an execution of $\mathcal{M}$.*

Proof. Observe that any two transactions in the execution fragment $E_1 \cdots E_m$ access mutually disjoint data sets. Since $\mathcal{M}$ is progressive and provides OF TM-liveness, there exists a t-sequential execution fragment $E' = E'_1 \cdots E'_m$ such that, for all $i \in \{1, \ldots, m\}$, the execution fragments E_i and E'_i are similar and $E \cdot E'$ is an execution of $\mathcal{M}$. Corollary 7.7 implies that, for all for all $i \in \{1, \ldots, m\}$, $\mathcal{M}$ has an execution of the form $E \cdot E'_1 \cdots E'_i \cdots E'_m \cdot \gamma_i$. More specifically, $\mathcal{M}$ has an execution of the form $E \cdot \gamma_i \cdot E'_1 \cdots E'_i \cdots E'_m$. Recall that the execution fragment ρ_i of fast-path transaction T_i that extends γ_i contains accesses only to base objects in $\bigcup_{X \in DSet(T_i)} \mathbb{D}_X$. Moreover, for all $i, j \in \{1, \ldots, m\}$; $i \neq j$, $Dset(T_i) \cap Dset(T_{m+j}) = \emptyset$ and $Dset(T_{m+i}) \cap Dset(T_{m+j}) = \emptyset$.

It follows that $\mathcal{M}$ has an execution of the form $E \cdot \gamma_i \cdot E'_1 \cdots E'_i \cdot \rho_i \cdot E'_{i+1} \cdots E'_m$. and the states of each of the base objects $\bigcup_{X \in DSet(T_i)} \mathbb{D}_X$ accessed by T_i in the configuration after $E \cdot \gamma_i \cdot E'_1 \cdots E'_i$ and $E \cdot \gamma_i \cdot E_i$ are the same. But $E \cdot \gamma_i \cdot E_i \cdot \rho_i$ is an execution of $\mathcal{M}$. Thus, for all $i \in \{1, \ldots, m\}$, $\mathcal{M}$ has an execution of the form $E \cdot E' \cdot \gamma_i \cdot \rho_i$. □

Finally, we are now ready to derive our lower bound.

Theorem 7.9. *Let $\mathcal{M}$ be any progressive, opaque HyTM implementation that provides OF TM-liveness. For every $m \in \mathbb{N}$, there exists an execution E in which some fast-path read-only transaction $T_k \in txns(E)$ satisfies either (1) $Dset(T_k) \leq m$ and T_k incurs a capacity abort in E or (2) $Dset(T_k) = m$ and T_k accesses $\Omega(m)$ distinct metadata base objects in E.*

Here is a high-level overview of the proof technique. Let κ be the smallest integer such that some fast-path transaction running step contention-free after a t-quiescent configuration performs κ t-reads and incurs a capacity abort.

We prove that, for all $m \leq \kappa - 1$, there exists a t-complete execution E_m and a set S_m with $|S_m| = 2^{\kappa - m}$ of read-only fast-path transactions that access mutually disjoint data sets such that each transaction in S_m that runs step contention-free from E_m and performs t-reads of m distinct t-objects accesses at least one distinct metadata base object within the execution of each t-read operation.

We proceed by induction. Assume that the induction statement holds for all $m < kappa - 1$. We prove that a set S_{m+1}; $|S_{m+1}| = 2^{\kappa-(m+1)}$ of fast-path transactions, each of which run step contention-free after the same t-complete execution E_{m+1}, perform $m + 1$ t-reads of distinct t-objects so that at least one distinct metadata base object is accessed within the execution of each t-read operation. In our construction, we pick any two new transactions from the set S_m and show that one of them running

step contention-free from a t-complete execution that extends E_m performs $m+1$ t-reads of distinct t-objects so that at least one distinct metadata base object is accessed within the execution of each t-read operation. In this way, the set of transactions is reduced by half in each step of the induction until one transaction remains which must have accessed a distinct metadata base object in every one of its $m+1$ t-reads.

Intuitively, since all the transactions that we use in our construction access mutually disjoint data sets, we can apply Lemma 7.6 to construct a t-complete execution E_{m+1} such that each of the fast-path transactions in S_{m+1} when running step contention-free after E_{m+1} perform $m+1$ t-reads so that at least one distinct metadata base object is accessed within the execution of each t-read operation.

We now present the formal proof:

Proof. In the constructions which follow, every fast-path transaction executes at most $m+1$ t-reads. Let κ be the smallest integer such that some fast-path transaction running step contention-free after a t-quiescent configuration performs κ t-reads and incurs a capacity abort. We proceed by induction.

Induction statement. We prove that, for all $m \leq \kappa - 1$, there exists a t-complete execution E_m and a set S_m with $|S_m| = 2^{\kappa - m}$ of read-only fast-path transactions that access mutually disjoint data sets such that each transaction $T_{f_i} \in S_m$ that runs step contention-free from E_m and performs t-reads of m distinct t-objects accesses at least one distinct metadata base object within the execution of each t-read operation. Let E_{f_i} be the step contention-free execution of T_{f_i} after E_m and let $\mathit{Dset}(T_{f_i}) = \{X_{i,1}, \ldots, X_{i,m}\}$.

The induction. Assume that the induction statement holds for all $m \leq \kappa - 1$. The statement is trivially true for the base case $m = 0$ for every $\kappa \in \mathbb{N}$.

We will prove that a set S_{m+1}; $|S_{m+1}| = 2^{\kappa-(m+1)}$ of fast-path transactions, each of which run step contention-free from the same t-quiescent configuration E_{m+1}, perform $m+1$ t-reads of distinct t-objects so that at least one distinct metadata base object is accessed within the execution of each t-read operation.

The construction proceeds in *phases*: there are exactly $\frac{|S_m|}{2}$ phases. In each phase, we pick any two new transactions from the set S_m and show that one of them running step contention-free after a t-complete execution that extends E_m performs $m+1$ t-reads of distinct t-objects so that at least one distinct metadata base object is accessed within the execution of each t-read operation.

Throughout this proof, we will assume that any two transactions (and resp. execution fragments) with distinct subscripts represent distinct identifiers.

For all $i \in \{0, \ldots, \frac{|S_m|}{2} - 1\}$, let $X_{2i+1}, X_{2i+2} \notin \bigcup_{i=0}^{|S_m|-1} \{X_{i,1}, \ldots, X_{i,m}\}$ be distinct t-objects and let v be the value of X_{2i+1} and X_{2i+2} after E_m. Let T_{s_i} denote a slow-path transaction which writes $nv \neq v$ to X_{2i+1} and X_{2i+2}. Let E_{s_i} be the t-complete step contention-free execution fragment of T_{s_i} running immediately after E_m.

Let E'_{s_i} be the longest prefix of the execution E_{s_i} such that $E_m \cdot E'_{s_i}$ can be extended neither with the complete step contention-free execution fragment of transaction $T_{f_{2i+1}}$ that performs its m t-reads of $X_{2i+1,1}, \ldots, X_{2i+1,m}$ and then performs $\mathit{read}_{f_{2i+1}}(X_{2i+1})$ and returns nv, nor with the complete step contention-free execution fragment of some transaction $T_{f_{2i+2}}$ that performs t-reads of $X_{2i+2_1}, \ldots, X_{2i+2,m}$ and then performs $\mathit{read}_{f_{2i+2}}(X_{2i+2})$ and returns nv. Progressiveness and OF TM-liveness of $\mathcal{M}$ stipulates that such an execution exists.

Let e_i be the enabled event of T_{s_i} in the configuration after $E_m \cdot E'_{s_i}$. By construction, the execution $E_m \cdot E'_{s_i}$ can be extended with at least one of the complete step contention-free executions of transaction $T_{f_{2i+1}}$ performing $(m+1)$ t-reads of $X_{2i+1,1}, \ldots, X_{2i+1,m}, X_{2i+1}$ such that $\mathit{read}_{f_{2i+1}}(X_{2i+1}) \to nv$ or transaction $T_{f_{2i+2}}$ performing t-reads of $X_{2i+2,1}, \ldots, X_{2i+2,m}, X_{2i+2}$ such that $\mathit{read}_{f_{2i+2}}(X_{2i+2}) \to nv$. Without loss of generality, suppose that $T_{f_{2i+1}}$ reads the value of X_{2i+1} to be nv after $E_m \cdot E'_{0_i} \cdot e_i$.

For any $i \in \{0, \ldots, \frac{|S_m|}{2} - 1\}$, we will denote by α_i the execution fragment which we will construct in phase i. For any $i \in \{0, \ldots, \frac{|S_m|}{2} - 1\}$, we prove that $\mathcal{M}$ has an execution of the form $E_m \cdot \alpha_i$ in which $T_{f_{2i+1}}$ (or $T_{f_{2i+2}}$) running step contention-free after a t-complete execution that extends E_m performs $m+1$ t-reads of distinct t-objects so that at least one distinct metadata base object is accessed within the execution of each first m t-read operations and $T_{f_{2i+1}}$ (or $T_{f_{2i+2}}$) is poised to apply an event after $E_m \cdot \alpha_i$ that accesses a distinct metadata base object during the $(m+1)^{th}$ t-read. Furthermore, we will show that $E_m \cdot \alpha_i$ appears t-sequential to $T_{f_{2i+1}}$ (or $T_{f_{2i+2}}$).

(Construction of phase i)

Let $E_{f_{2i+1}}$ (and resp. $E_{f_{2i+2}}$) be the complete step contention-free execution of the t-reads of $X_{2i+1,1}, \ldots, X_{2i+1,m}$ (and resp. $X_{2i+2,1}, \ldots, X_{2i+2,m}$) running after E_m by $T_{f_{2i+1}}$ (and resp. $T_{f_{2i+2}}$). By the inductive hypothesis, transaction $T_{f_{2i+1}}$ (and resp. $T_{f_{2i+2}}$) accesses m distinct metadata objects in the execution $E_m \cdot E_{f_{2i+1}}$ (and resp. $E_m \cdot E_{f_{2i+2}}$). Recall that transaction $T_{f_{2i+1}}$ does not conflict with transaction T_{s_i}. Thus, by Corollary 7.7, $\mathcal{M}$ has an execution of the form $E_m \cdot E'_{s_i} \cdot e_i \cdot E_{f_{2i+1}}$ (and resp. $E_m \cdot E'_{s_i} \cdot e_i \cdot E_{f_{2i+2}}$).

Let $E_{rf_{2i+1}}$ be the complete step contention-free execution fragment of $read_{f_{2i+1}}(X_{2i+1})$ that extends $E_{2i+1} = E_m \cdot E'_{s_i} \cdot e_i \cdot E_{f_{2i+1}}$. By OF TM-liveness, $read_{f_{2i+1}}(X_{2i+1})$ must return a matching response in $E_{2i+1} \cdot E_{rf_{2i+1}}$. We now consider two cases.

***Case I**: Suppose $E_{rf_{2i+1}}$ accesses at least one metadata base object b not previously accessed by $T_{f_{2i+1}}$.*

Let $E'_{rf_{2i+1}}$ be the longest prefix of $E_{rf_{2i+1}}$ which does not apply any primitives to any metadata base object b not previously accessed by $T_{f_{2i+1}}$. The execution $E_m \cdot E'_{s_i} \cdot e_i \cdot E_{f_{2i+1}} \cdot E'_{rf_{2i+1}}$ appears t-sequential to $T_{f_{2i+1}}$ because $E_{f_{2i+1}}$ does not contend with T_{s_i} on any base object and any common base object accessed in the execution fragments $E'_{rx_{2i+1}}$ and E_{s_i} by $T_{f_{2i+1}}$ and T_{s_i} respectively must be data objects contained in $\mathbb{D}$. Thus, we have that $|Dset(T_{f_{2i+1}})| = m+1$ and that $T_{f_{2i+1}}$ accesses m distinct metadata base objects within each of its first m t-read operations and is poised to access a distinct metadata base object during the execution of the $(m+1)^{th}$ t-read. In this case, let $\alpha_i = E_m \cdot E'_{s_i} \cdot e_i \cdot E_{f_{2i+1}} \cdot E'_{rf_{2i+1}}$.

***Case II**: Suppose $E_{rf_{2i+1}}$ does not access any metadata base object not previously accessed by $T_{f_{2i+1}}$.*

In this case, we will first prove the following:

Claim 7.10. *$\mathcal{M}$ has an execution of the form $E_{2i+2} = E_m \cdot E'_{s_i} \cdot e_i \cdot \bar{E}_{f_{2i+1}} \cdot E_{f_{2i+2}}$ where $\bar{E}_{f_{2i+1}}$ is the t-complete step contention-free execution of $T_{f_{2i+1}}$ in which $read_{f_{2i+1}}(X_{2i+1}) \to nv$, $T_{f_{2i+1}}$ invokes $tryC_{f_{2i+1}}$ and returns a matching response.*

Proof. Since $E_{rf_{2i+1}}$ does not contain accesses to any distinct metadata base objects, the execution $E_m \cdot E'_{s_i} \cdot e_i \cdot E_{f_{2i+1}} \cdot E_{rf_{2i+1}}$ appears t-sequential to $T_{f_{2i+1}}$. By definition of the event e_i, $read_{f_{2i+1}}(X_{2i+1})$ must access the base object to which the event e_i applies a nontrivial primitive and return the response nv in $E'_{s_i} \cdot e_i \cdot E_{f_{2i+1}} \cdot E_{rf_{2i+1}}$. By OF TM-liveness, it follows that $E_m \cdot E'_{s_i} \cdot e_i \cdot \bar{E}_{f_{2i+1}}$ is an execution of $\mathcal{M}$.

Now recall that $E_m \cdot E'_{s_i} \cdot e_i \cdot E_{f_{2i+2}}$ is an execution of $\mathcal{M}$ because transactions $T_{f_{2i+2}}$ and T_{s_i} do not conflict in this execution and thus, cannot contend on any base object. Finally, because $T_{f_{2i+1}}$ and $T_{f_{2i+2}}$ access disjoint data sets in $E_m \cdot E'_{s_i} \cdot e_i \cdot \bar{E}_{f_{2i+1}} \cdot E_{f_{2i+2}}$, by Lemma 7.6 again, we have that $E_m \cdot E'_{s_i} \cdot e_i \cdot \bar{E}_{f_{2i+1}} \cdot E_{f_{2i+2}}$ is an execution of $\mathcal{M}$. □

Let $E_{rf_{2i+2}}$ be the complete step contention-free execution fragment of $read_{f_{2i+2}}(X_{2i+2})$ after $E_m \cdot E'_{s_i} \cdot e_i \cdot \bar{E}_{f_{2i+1}} \cdot E_{f_{2i+2}}$. By the induction hypothesis and Claim 7.10, transaction $T_{f_{2i+2}}$ must access m distinct metadata base objects in the execution $E_m \cdot E'_{s_i} \cdot e_i \cdot \bar{E}_{f_{2i+1}} \cdot E_{f_{2i+2}}$.

If $E_{rf_{2i+2}}$ accesses some metadata base object, then by the argument given in Case I applied to transaction $T_{f_{2i+2}}$, we get that $T_{f_{2i+2}}$ accesses m distinct metadata base objects within each of the first m t-read operations and is poised to access a distinct metadata base object during the execution of the $(m+1)^{th}$ t-read.

Thus, suppose that $E_{rf_{2i+2}}$ does not access any metadata base object previously accessed by $T_{f_{2i+2}}$. We claim that this is impossible and proceed to derive a contradiction. In particular, $E_{rf_{2i+2}}$ does not contend with T_{s_i} on any metadata base object. Consequently, the execution $E_m \cdot E'_{s_i} \cdot e_i \cdot \bar{E}_{f_{2i+1}} \cdot E_{f_{2i+2}}$ appears t-sequential to $T_{x_{2i+2}}$ since $E_{rx_{2i+2}}$ only contends with T_{s_i} on base objects in $\mathbb{D}$. It follows that $E_{2i+2} \cdot E_{rf_{2i+2}}$ must also appear t-sequential to $T_{f_{2i+2}}$ and so $E_{rf_{2i+2}}$ cannot abort. Recall that the base object, say b, to which T_{s_i} applies a nontrivial primitive in the event e_i is accessed by $T_{f_{2i+1}}$ in $E_m \cdot E'_{s_i} \cdot e_i \cdot \bar{E}_{f_{2i+1}} \cdot E_{f_{2i+2}}$; thus, $b \in \mathbb{D}_{X_{2i+1}}$. Since $X_{2i+1} \notin \mathit{Dset}(T_{f_{2i+2}})$, b cannot be accessed by $T_{f_{2i+2}}$. Thus, the execution $E_m \cdot E'_{s_i} \cdot e_i \cdot \bar{E}_{f_{2i+1}} \cdot E_{f_{2i+2}} \cdot E_{rf_{2i+2}}$ is indistinguishable to $T_{f_{2i+2}}$ from the execution $\hat{E}_i \cdot E'_{s_i} \cdot E_{f_{2i+2}} \cdot E_{rf_{2i+2}}$ in which $\mathit{read}_{f_{2i+2}}(X_{2i+2})$ must return the response v (by construction of E'_{s_i}).

But we observe now that the execution $E_m \cdot E'_{s_i} \cdot e_i \cdot \bar{E}_{f_{2i+1}} \cdot E_{f_{2i+2}} \cdot E_{rf_{2i+2}}$ is not opaque. In any serialization corresponding to this execution, T_{s_i} must be committed and must precede $T_{f_{2i+1}}$ because $T_{f_{2i+1}}$ read nv from X_{2i+1}. Also, transaction $T_{f_{2i+2}}$ must precede T_{s_i} because $T_{f_{2i+2}}$ read v from X_{2i+2}. However $T_{f_{2i+1}}$ must precede $T_{f_{2i+2}}$ to respect real-time ordering of transactions. Clearly, there exists no such serialization—contradiction.

Letting $E'_{rf_{2i+2}}$ be the longest prefix of $E_{rf_{2i+2}}$ which does not access a base object $b \in \mathbb{M}$ not previously accessed by $T_{f_{2i+2}}$, we can let $\alpha_i = E'_{s_i} \cdot e_i \cdot \bar{E}_{f_{2i+1}} \cdot E_{f_{2i+2}} \cdot E'_{rf_{2i+2}}$ in this case.

Combining Cases I and II, the following claim holds.

Claim 7.11. *For each $i \in \{0, \ldots, \frac{|S_m|}{2} - 1\}$, $\mathcal{M}$ has an execution of the form $E_m \cdot \alpha_i$ in which*

(1) some fast-path transaction $T_i \in \mathit{txns}(\alpha_i)$ performs t-reads of $m+1$ distinct t-objects so that at least one distinct metadata base object is accessed within the execution of each of the first m t-reads, T_i is poised to access a distinct metadata base object after $E_m \cdot \alpha_i$ during the execution of the $(m+1)^{th}$ t-read and the execution appears t-sequential to T_i,

(2) the two fast-path transactions in the execution fragment α_i do not contend on the same base object.

(Collecting the phases)

We will now describe how we can construct the set S_{m+1} of fast-path transactions from these $\frac{|S_m|}{2}$ phases and force each of them to access $m+1$ distinct metadata base objects when running step contention-free after the same t-complete execution.

For each $i \in \{0, \ldots, \frac{|S_m|}{2} - 1\}$, let β_i be the subsequence of the execution α_i consisting of all the events of the fast-path transaction that is poised to access a $(m+1)^{th}$ distinct metadata base object. Henceforth, we denote by T_i the fast-path transaction that participates in β_i. Then, from Claim 7.11, it follows that, for each $i \in \{0, \ldots, \frac{|S_m|}{2} - 1\}$, $\mathcal{M}$ has an execution of the form $E_m \cdot E'_{s_i} \cdot e_i \cdot \beta_i$ in which the fast-path transaction T_i performs t-reads of $m+1$ distinct t-objects so that at least one distinct metadata base object is accessed within the execution of each of the first m t-reads, T_i is poised to access a distinct metadata base object after $E_m \cdot E'_{s_i} \cdot e_i \cdot \beta_i$ during the execution of the $(m+1)^{th}$ t-read and the execution appears t-sequential to T_i.

The following result is a corollary to the above claim that is obtained by applying the definition of "appears t-sequential". Recall that $E'_{s_i} \cdot e_i$ is the t-incomplete execution of slow-path transaction T_{s_i} that accesses t-objects X_{2i+1} and X_{2i+2}.

Corollary 7.12. *For all $i \in \{0, \ldots, \frac{|(S_m|}{2} - 1\}$, $\mathcal{M}$ has an execution of the form $E_m \cdot E_i \cdot \beta_i$ such that the configuration after $E_m \cdot E_i$ is t-quiescent, $\mathit{txns}(E_i) \subseteq \{T_{s_i}\}$ and $\mathit{Dset}(T_{s_i}) \subseteq \{X_{2i+1}, X_{2i+2}\}$ in E_i.*

We can represent the execution $\beta_i = \gamma_i \cdot \rho_i$ where fast-path transaction T_i performs complete t-reads of m distinct t-objects in γ_i and then performs an incomplete t-read of the $(m+1)^{th}$ t-object in ρ_i in which T_i only accesses base objects in $\bigcup_{X \in DSet(T_i)} \{X\}$. Recall that T_i and T_{s_i} do not contend on the same base object in the execution $E_m \cdot E_i \cdot \gamma_i$. Thus, for all $i \in \{0, \ldots, \frac{|S_m|}{2} - 1\}$, $\mathcal{M}$ has an execution of the form $E_m \cdot \gamma_i \cdot E_i \cdot \rho_i$.

Observe that the fast-path transaction $T_i \in \gamma_i$ does not access any t-object that is accessed by any slow-path transaction in the execution fragment $E_0 \cdots E_{\frac{|S_m|}{2}-1}$. By Lemma 7.8, there exists a t-complete step contention-free execution fragment E' that is similar to $E_0 \cdots E_{\frac{|S_m|}{2}-1}$ such that for all $i \in \{0, \ldots, \frac{|S_m|}{2} - 1\}$, $\mathcal{M}$ has an execution of the form $E_m \cdot E' \cdot \gamma_i \cdot \rho_i$. By our construction, the enabled event of each fast-path transaction $T_i \in \beta_i$ in this execution is an access to a distinct metadata base object.

Let S_{m+1} denote the set of all fast-path transactions that participate in the execution fragment $\beta_0 \cdots \beta_{\frac{|(S_m|}{2}-1}$ and $E_{m+1} = E_m \cdot E'$. Thus, $|S_{m+1}|$ fast-path transactions, each of which run step contention-free from the same t-quiescent configuration, perform $m + 1$ t-reads of distinct t-objects so that at least one distinct metadata base object is accessed within the execution of each t-read operation. This completes the proof. □

7.4 Instrumentation-optimal progressive HyTM

We prove that the lower bound in Theorem 7.9 is tight by describing an 'instrumentation-optimal" HyTM implementation (Algorithm 7.1) that is opaque, progressive, provides wait-free TM-liveness, uses *invisible reads*.

Base objects. For every t-object X_j, our implementation maintains a base object $v_j \in \mathbb{D}$ that stores the value of X_j and a metadata base object r_j, which is a *lock bit* that stores 0 or 1.

Fast-path transactions. For a fast-path transaction T_k, the $read_k(X_j)$ implementation first reads r_j to check if X_j is locked by a concurrent updating transaction. If so, it returns A_k, else it returns the value of X_j. Updating fast-path transactions use uninstrumented writes: $write(X_j, v)$ simply stores the cached state of X_j along with its value v and if the cache has not been invalidated, updates the shared memory during $tryC_k$ by invoking the *commit-cache* primitive.

Slow-path read-only transactions. Any $read_k(X_j)$ invoked by a slow-path transaction first reads the value of the object from v_j, checks if r_j is set and then performs *value-based validation* on its entire read set to check if any of them have been modified. If either of these conditions is true, the transaction returns A_k. Otherwise, it returns the value of X_j. A read-only transaction simply returns C_k during the tryCommit.

Slow-path updating transactions. The $write_k(X, v)$ implementation of a slow-path transaction stores v and the current value of X_j locally, deferring the actual update in shared memory to tryCommit.

During $tryC_k$, an updating slow-path transaction T_k attempts to obtain exclusive write access to its entire write set as follows: for every t-object $X_j \in Wset(T_k)$, it writes 1 to each base object r_j by performing a *compare-and-set* (*cas*) primitive that checks if the value of r_j is not 1 and, if so, replaces it with 1. If the *cas* fails, then T_k releases the locks on all objects X_ℓ it had previously acquired by writing 0 to r_ℓ and then returns A_k. Intuitively, if the *cas* fails, some concurrent transaction is performing a t-write to a t-object in $Wset(T_k)$. If all the locks on the write set were acquired successfully, T_k checks if any t-object in $Rset(T_k)$ is concurrently being updated by another transaction and then performs value-based validation of the read set. If a conflict is detected from the these checks, the transaction is aborted. Finally, $tryC_k$ attempts to write the values of the t-objects via *cas* operations. If any *cas* on the individual base objects fails, there must be a concurrent fast-path writer, and so T_k rolls back the state of the base objects that were updated, releases locks on its write set and returns A_k. The roll backs are performed with *cas* operations, skipping any which fail to allow for concurrent fast-path writes to locked locations. Note that if a concurrent read operation of a fast-path transaction T_ℓ finds an "invalid" value in v_j that was written by such transaction T_k but has not been rolled back yet, then T_ℓ either incurs a tracking set abort later because T_k has updated v_j or finds r_j to be 1. In both cases, the read operation of T_ℓ aborts.

The implementation uses invisible reads (no nontrivial primitives are applied by reading transactions). Every t-operation returns a matching response within a finite number of its steps.

Algorithm 7.1 Progressive opaque HyTM implementation that provides uninstrumented writes and invisible reads; code for process p_i executing transaction T_k

```
1:  Shared objects:
2:    v_j ∈ 𝔻, for each t-object X_j
3:        allows reads, writes and cas
4:    r_j ∈ 𝕄, for each t-object X_j
5:        allows reads, writes and cas

6:  Local objects:
7:    Lset(T_k) ⊆ Wset(T_k), initially empty
8:    Oset(T_k) ⊆ Wset(T_k), initially empty
      Code for slow-path transactions

9:  read_k(X_j):   // slow-path
10:   if X_j ∉ Rset_k then
11:     [ov_j, k_j] := read(v_j)
12:     Rset(T_k) := Rset(T_k) ∪ {X_j, [ov_j, k_j]}
13:     if r_j ≠ 0 then
14:       Return A_k
15:     if ∃X_j ∈ Rset(T_k):(ov_j, k_j) ≠ read(v_j) then
16:       Return A_k
17:     Return ov_j
18:   else
19:     ov_j := Rset(T_k).locate(X_j)
20:     Return ov_j

21: write_k(X_j, v):   // slow-path
22:   (ov_j, k_j) := read(v_j)
23:   nv_j := v
24:   Wset(T_k) := Wset(T_k) ∪ {X_j, [ov_j, k_j]}
25:   Return ok

26: tryC_k():   // slow-path
27:   if Wset(T_k) = ∅ then
28:     Return C_k
29:   locked := acquire(Wset(T_k))
30:   if ¬ locked then
31:     Return A_k
32:   if isAbortable() then
33:     release(Lset(T_k))
34:     Return A_k
35:   for all X_j ∈ Wset(T_k) do
36:     if v_j.cas([ov_j, k_j], [nv_j, k]) then
37:       Oset(T_k) := Oset(T_k) ∪ {X_j}
38:     else
39:       undo(Oset(T_k))
40:   release(Wset(T_k))
41:   Return C_k

42: Function: acquire(Q):
43:   for all X_j ∈ Q do
44:     if r_j.cas(0, 1) then
45:       Lset(T_k) := Lset(T_k) ∪ {X_j}
46:     else
47:       release(Lset(T_k))
48:       Return false
49:   Return true

50: Function: release(Q):
51:   for all X_j ∈ Q do
52:     r_j.write(0)
53:   Return ok

54: Function: undo(Oset(T_k)):
55:   for all X_j ∈ Oset(T_k) do
56:     v_j.cas([nv_j, k], [ov_j, k_j])
57:   release(Wset(T_k))
58:   Return A_k

59: Function: isAbortable() :
60:   if ∃X_j ∈ Rset(T_k): X_j ∉ Wset(T_k) ∧ read(r_j) ≠ 0 then
61:     Return true
62:   if ∃X_j ∈ Rset(T_k):[ov_j, k_j] ≠ read(v_j) then
63:     Return true
64:   Return false

      Code for fast-path transactions

65: read_k(X_j):   // fast-path
66:   [ov_j, k_j] := read(v_j) // cached read
67:   if read(r_j) ≠ 0 then
68:     Return A_k
69:   Return ov_j

70: write_k(X_j, v):   // fast-path
71:   write(v_j, [nv_j, k]) // cached write
72:   Return ok

73: tryC_k():   // fast-path
74:   commit-cache_i // returns C_k or A_k
```

Complexity. Every t-read operation performed by a fast-path transaction accesses a metadata base object once (the lock bit corresponding to the t-object), which is the price to pay for detecting conflicting updating slow-path transactions. Write operations of fast-path transactions are uninstrumented.

Lemma 7.13. *Algorithm 7.1 implements an opaque TM.*

Proof. Let E by any execution of Algorithm 7.1. Since opacity is a safety property, it is sufficient to prove that every finite execution is opaque [18]. Let $<_E$ denote a total-order on events in E.

Let H denote a subsequence of E constructed by selecting *linearization points* of t-operations performed in E. The linearization point of a t-operation *op*, denoted as ℓ_{op} is associated with a base object event or an event performed during the execution of *op* using the following procedure.

Completions. First, we obtain a completion of E by removing some pending invocations or adding responses to the remaining pending invocations as follows:

- incomplete $read_k$, $write_k$ operation performed by a slow-path transaction T_k is removed from E; an incomplete $tryC_k$ is removed from E if T_k has not performed any write to a base object r_j; $X_j \in Wset(T_k)$ in Line 36, otherwise it is completed by including C_k after E.
- every incomplete $read_k$, $tryA_k$, $write_k$ and $tryC_k$ performed by a fast-path transaction T_k is removed from E.

Linearization points. Now a linearization H of E is obtained by associating linearization points to t-operations in the obtained completion of E. For all t-operations performed a slow-path transaction T_k, linearization points as assigned as follows:

- For every t-read op_k that returns a non-A_k value, ℓ_{op_k} is chosen as the event in Line 11 of Algorithm 7.1, else, ℓ_{op_k} is chosen as invocation event of op_k
- For every $op_k = write_k$ that returns, ℓ_{op_k} is chosen as the invocation event of op_k
- For every $op_k = tryC_k$ that returns C_k such that $Wset(T_k) \neq \emptyset$, ℓ_{op_k} is associated with the first write to a base object performed by release when invoked in Line 40, else if op_k returns A_k, ℓ_{op_k} is associated with the invocation event of op_k
- For every $op_k = tryC_k$ that returns C_k such that $Wset(T_k) = \emptyset$, ℓ_{op_k} is associated with Line 28

For all t-operations performed a fast-path transaction T_k, linearization points as assigned as follows:

- For every t-read op_k that returns a non-A_k value, ℓ_{op_k} is chosen as the event in Line 66 of Algorithm 7.1, else, ℓ_{op_k} is chosen as invocation event of op_k
- For every op_k that is a $tryC_k$, ℓ_{op_k} is the $commit\text{-}cache_k$ primitive invoked by T_k
- For every op_k that is a $write_k$, ℓ_{op_k} is the event in Line 71.

$<_H$ denotes a total-order on t-operations in the complete sequential history H.

Serialization points. The serialization of a transaction T_j, denoted as δ_{T_j} is associated with the linearization point of a t-operation performed by the transaction.

We obtain a t-complete history $\bar{H}$ from H as follows. A serialization S is obtained by associating serialization points to transactions in $\bar{H}$ as follows: for every transaction T_k in H that is complete, but not t-complete, we insert $tryC_k \cdot A_k$ immediately after the last event of T_k in H.

- If T_k is an updating transaction that commits, then δ_{T_k} is ℓ_{tryC_k}
- If T_k is a read-only or aborted transaction, then δ_{T_k} is assigned to the linearization point of the last t-read that returned a non-A_k value in T_k

$<_S$ denotes a total-order on transactions in the t-sequential history S.

Claim 7.14. *If $T_i \prec_H T_j$, then $T_i <_S T_j$*

Proof. This follows from the fact that for a given transaction, its serialization point is chosen between the first and last event of the transaction implying if $T_i \prec_H T_j$, then $\delta_{T_i} <_E \delta_{T_j}$ implies $T_i <_S T_j$. □

Claim 7.15. *S is legal.*

Proof. We claim that for every $read_j(X_m) \to v$, there exists some slow-path transaction T_i (or resp. fast-path) that performs $write_i(X_m, v)$ and completes the event in Line 36 (or resp. Line 71) such that $read_j(X_m) \not\prec_H^{RT} write_i(X_m, v)$.

Suppose that T_i is a slow-path transaction: since $read_j(X_m)$ returns the response v, the event in Line 11 succeeds the event in Line 36 performed by $tryC_i$. Since $read_j(X_m)$ can return a non-abort response only after T_i writes 0 to r_m in Line 52, T_i must be committed in S. Consequently, $\ell_{tryC_i} <_E \ell_{read_j(X_m)}$. Since, for any updating committing transaction T_i, $\delta_{T_i} = \ell_{tryC_i}$, it follows that $\delta_{T_i} <_E \delta_{T_j}$.

Otherwise if T_i is a fast-path transaction, then clearly T_i is a committed transaction in S. Recall that $read_j(X_m)$ can read v during the event in Line 11 only after T_i applies the *commit-cache* primitive. By the assignment of linearization points, $\ell_{tryC_i} <_E \ell_{read_j(X_m)}$ and thus, $\delta_{T_i} <_E \ell_{read_j(X_m)}$.

Thus, to prove that S is legal, it suffices to show that there does not exist a transaction T_k that returns C_k in S and performs $write_k(X_m, v')$; $v' \neq v$ such that $T_i <_S T_k <_S T_j$.

T_i and T_k are both updating transactions that commit. Thus,

$$(T_i <_S T_k) \iff (\delta_{T_i} <_E \delta_{T_k})$$
$$(\delta_{T_i} <_E \delta_{T_k}) \iff (\ell_{tryC_i} <_E \ell_{tryC_k})$$

Since, T_j reads the value of X written by T_i, one of the following is true: $\ell_{tryC_i} <_E \ell_{tryC_k} <_E \ell_{read_j(X_m)}$ or $\ell_{tryC_i} <_E \ell_{read_j(X_m)} <_E \ell_{tryC_k}$.

Suppose that $\ell_{tryC_i} <_E \ell_{tryC_k} <_E \ell_{read_j(X_m)}$.

(*Case I:*) T_i and T_k are slow-path transactions.

Thus, T_k returns a response from the event in Line 29 before the read of the base object associated with X_m by T_j in Line 11. Since T_i and T_k are both committed in E, T_k returns *true* from the event in Line 29 only after T_i writes 0 to r_m in Line 52.

If T_j is a slow-path transaction, recall that $read_j(X_m)$ checks if X_j is locked by a concurrent transaction, then performs read-validation (Line 13) before returning a matching response. We claim that $read_j(X_m)$ must return A_j in any such execution.

Consider the following possible sequence of events: T_k returns *true* from *acquire* function invocation, updates the value of X_m to shared-memory (Line 36), T_j reads the base object v_m associated with X_m, T_k releases X_m by writing 0 to r_m and finally T_j performs the check in Line 13. But in this case, $read_j(X_m)$ is forced to return the value v' written by T_m— contradiction to the assumption that $read_j(X_m)$ returns v.

Otherwise suppose that T_k acquires exclusive access to X_m by writing 1 to r_m and returns *true* from the invocation of *acquire*, updates v_m in Line 36), T_j reads v_m, T_j performs the check in Line 13 and finally T_k releases X_m by writing 0 to r_m. Again, $read_j(X_m)$ must return A_j since T_j reads that r_m is 1—contradiction.

A similar argument applies to the case that T_j is a fast-path transaction. Indeed, since every *data* base object read by T_j is contained in its tracking set, if any concurrent transaction updates any t-object in its read set, T_j is aborted immediately by our model(cf. Section 7.2.2).

Thus, $\ell_{tryC_i} <_E \ell_{read_j(X)} <_E \ell_{tryC_k}$.

(*Case II:*) T_i is a slow-path transaction and T_k is a fast-path transaction. Thus, T_k returns C_k before the read of the base object associated with X_m by T_j in Line 11, but after the response of *acquire* by T_i in Line 29. Since $read_j(X_m)$ reads the value of X_m to be v and not v', T_i performs the *cas* to v_m in Line 36 after the T_k performs the *commit-cache* primitive (since if otherwise, T_k would be aborted in E). But then the *cas* on v_m performed by T_i would return *false* and T_i would return A_i—contradiction.

(*Case III:*) T_k is a slow-path transaction and T_i is a fast-path transaction. This is analogous to the above case.

(*Case IV:*) T_i and T_k are fast-path transactions. Thus, T_k returns C_k before the read of the base object associated with X_m by T_j in Line 11, but before T_i returns C_i (this follows from Observations 7.1 and 7.2). Consequently, $read_j(X_m)$ must read the value of X_m to be v' and return v'—contradiction.

We now need to prove that δ_{T_j} indeed precedes ℓ_{tryC_k} in E.

Consider the two possible cases:

- Suppose that T_j is a read-only transaction. Then, δ_{T_j} is assigned to the last t-read performed by T_j that returns a non-A$_j$ value. If $read_j(X_m)$ is not the last t-read that returned a non-A$_j$ value, then there exists a $read_j(X')$ such that $\ell_{read_j(X_m)} <_E \ell_{tryC_k} <_E \ell_{read_j(X')}$. But then this t-read of X' must abort by performing the checks in Line 13 or incur a tracking set abort—contradiction.
- Suppose that T_j is an updating transaction that commits, then $\delta_{T_j} = \ell_{tryC_j}$ which implies that $\ell_{read_j(X)} <_E \ell_{tryC_k} <_E \ell_{tryC_j}$. Then, T_j must neccesarily perform the checks in Line 32 and return A_j or incur a tracking set abort—contradiction to the assumption that T_j is a committed transaction.

The proof follows. □

The conjunction of Claims 7.14 and 7.15 establish that Algorithm 7.1 is opaque. □

Theorem 7.16. *There exists an opaque HyTM implementation that provides uninstrumented writes, invisible reads, progressiveness and wait-free TM-liveness such that in its every execution E, every read-only fast-path transaction $T \in txns(E)$ accesses $O(|Rset(T)|)$ distinct metadata base objects.*

Proof. (Opacity) Follows from Lemma 7.13.

(TM-liveness and TM-progress) Since none of the implementations of the t-operations in Algorithm 7.1 contain unbounded loops or waiting statements, Algorithm 7.1 provides wait-free TM-liveness, *i.e.*, every t-operation returns a matching response after taking a finite number of steps.

Consider the cases under which a slow-path transaction T_k may be aborted in any execution.

- Suppose that there exists a $read_k(X_j)$ performed by T_k that returns A_k from Line 13. Thus, there exists a transaction that has written 1 to r_j in Line 44, but has not yet written 0 to r_j in Line 52 or some t-object in $Rset(T_k)$ has been updated since its t-read by T_k. In both cases, there exists a concurrent transaction performing a t-write to some t-object in $Rset(T_k)$, thus forcing a read-write conflict.
- Suppose that $tryC_k$ performed by T_k that returns A_k from Line 30. Thus, there exists a transaction that has written 1 to r_j in Line 44, but has not yet written 0 to r_j in Line 52. Thus, T_k encounters write-write conflict with another transaction that concurrently attempts to update a t-object in $Wset(T_k)$.
- Suppose that $tryC_k$ performed by T_k that returns A_k from Line 32. Since T_k returns A_k from Line 32 for the same reason it returns A_k after Line 13, the proof follows.

Consider the cases under which a fast-path transaction T_k may be aborted in any execution E.

- Suppose that a $read_k(X_m)$ performed by T_k returns A_k from Line 67. Thus, there exists a concurrent slow-path transaction that is pending in its tryCommit and has written 1 to r_m, but not released the lock on X_m i.e. T_k conflicts with another transaction in E.
- Suppose that T_k returns A_k while performing a cached access of some base object b via a trivial (and resp. nontrivial) primitive. Indeed, this is possible only if some concurrent transaction writes (and resp. reads or writes) to b. However, two transactions T_k and T_m may contend on b in E only if there exists $X \in Dset(T_i) \cap Dset(T_j)$ and $X \in Wset(T_i) \cup Wset(T_j)$. from Line 30. The same argument applies for the case when T_k returns A_k while performing *commit-cache*$_k$ in E.

(Complexity) The implementation uses uninstrumented writes since each $write_k(X_m)$ simply writes to $v_m \in \mathbb{D}_{X_m}$ and does not access any metadata base object. The complexity of each $read_k(X_m)$ is a single access to a metadata base object r_m in Line 67 that is not accessed any other transaction T_i unless $X_m \in Dset(T_i)$. while the $tryC_k$ just calls $cache\text{-}commit_k$ that returns C_k. Thus, each read-only transaction T_k accesses $O(|Rset(T_k)|)$ distinct metadata base objects in any execution. □

7.5 Providing partial concurrency at low cost

Algorithm 7.2 Opaque HyTM implementation with progressive slow-path and sequential fast-path TM-progress; code for T_k by process p_i

```
1:  Shared objects:
2:    v_j ∈ 𝔻, for each t-object X_j
3:        allows reads, writes and cas
4:    r_j ∈ 𝕄, for each t-object X_j
5:        allows reads, writes and cas
6:    fa, fetch-and-add object
      Code for slow-path transactions
7:  tryC_k(): // slow-path
8:    if Wset(T_k) = ∅ then
9:      Return C_k
10:   locked := acquire(Wset(T_k))
11:   if ¬ locked then
12:     Return A_k
13:   fa.add(1)
14:   if isAbortable() then
15:     release(Lset(T_k))
16:     Return A_k
17:   for all X_j ∈ Wset(T_k) do
18:     if v_j.cas((ov_j, k_j), (nv_j, k)) then
19:       Oset(T_k) := Oset(T_k) ∪ {X_j}
20:     else
21:       Return undo(Oset(T_k))
22:   release(Wset(T_k))
23:   Return C_k

24: Function: release(Q):
25:   for all X_j ∈ Q do
26:     r_j.write(0)
27:   fa.add(−1)
28:   Return ok

      Code for fast-path transactions
29: read_k(X_j): // fast-path
30:   if Rset(T_k) = ∅ then
31:     l ← read(fa) // cached read
32:     if l ≠ 0 then
33:       Return A_k
34:   (ov_j, k_j) := read(v_j) // cached read
35:   Return ov_j

36: write_k(X_j, v): // fast-path
37:   v_j.write(nv_j, k) // cached write
38:   Return ok

39: tryC_k(): // fast-path
40:   commit-cache_i // returns C_k or A_k
```

We showed that allowing fast-path transactions to run concurrently in HyTM results in an instrumentation cost that is proportional to the read-set size of a fast-path transaction. But can we run at least *some* transactions concurrently with constant instrumentation cost, while still keeping invisible reads?

Algorithm 7.2 implements a *slow-path progressive* opaque HyTM with invisible reads and wait-free TM-liveness. To fast-path transactions, it only provides *sequential* TM-progress (they are only guaranteed to commit in the absence of concurrency), but in return the algorithm is only using a single metadata base object *fa* that is read once by a fast-path transaction and accessed twice with a *fetch-and-add* primitive by an updating slow-path transaction. Thus, the instrumentation cost of the algorithm is constant.

Intuitively, *fa* allows fast-path transactions to detect the existence of concurrent updating slow-path transactions. Each time an updating slow-path updating transaction tries to commit, it increments *fa* and once all writes to data base objects are completed (this part of the algorithm is identical to Algorithm 7.1) or the transaction is aborted, it decrements *fa*. Therefore, $fa \neq 0$ means that at least one slow-path updating transaction is incomplete. A fast-path transaction simply checks if $fa \neq 0$ in the beginning and aborts if so, otherwise, its code is identical to that in Algorithm 7.1. Note that this way, any update of *fa* automatically causes a tracking set abort of any incomplete fast-path transaction.

Theorem 7.17. *There exists an opaque HyTM implementation that provides uninstrumented writes, invisible reads, progressiveness for slow-path transactions, sequential TM-progress for fast-path transactions and wait-free TM-liveness such that in every its execution E, every fast-path transaction accesses at most one metadata base object.*

Proof. The proof of opacity is almost identical to the analogous proof for Algorithm 7.1 in Lemma 7.13.

As with Algorithm 7.1, enumerating the cases under which a slow-path transaction T_k returns A_k proves that Algorithm 7.2 satisfies progressiveness for slow-path transactions. Any fast-path transaction T_k; $Rset(T_k) \neq \emptyset$ reads the metadata base object *fa* and adds it to the process's tracking set (Line 31). If the value of *fa* is not 0, indicating that there exists a concurrent slow-path transaction pending in its tryCommit, T_k returns A_k. Thus, the implementation provides sequential TM-progress for fast-path transactions.

Also, in every execution E of $\mathcal{M}$, no fast-path write-only transaction accesses any metadata base object and a fast-path reading transaction accesses the metadata base object *fa* exactly once, during the first t-read. □

7.6 Related work and Discussion

HyTM model. Our HyTM model is a natural extension of the model we specified for Software Transactional memory (cf. Chapter 2), and has the advantage of being relatively simple. The term *instrumentation* was originally used in the context of HyTMs [35, 99, 112] to indicate the overhead a hardware transaction induces in order to detect pending software transactions. The impossibility of designing HyTMs without any code instrumentation was intuitively suggested in [35], we present a formal proof in this paper.

In [23], Attiya and Hillel considered the instrumentation cost of *privatization*, *i.e.*, allowing transactions to isolate data items by making them private to a process so that no other process is allowed to modify the privatized item. Just as we capture a tradeoff between the cost of hardware instrumentation and the amount of concurrency allowed between hardware and software transactions, [23] captures a tradeoff between the cost of privatization and the number of transactions guaranteed to make progress concurrently in ℓ-progressive STMs. The model we consider is fundamentally different to [23], in that we model hardware transactions at the level of cache coherence, and do not consider non-transactional accesses, *i.e.*, neither data nor meta-data base objects are private in our HyTM model. The proof techniques we employ are also different.

Uninstrumented HTMs may be viewed as being *disjoint-access parallel (DAP)* [24, 86]. As such, some of the techniques used in the proof of Theorem 7.4 resemble those used in [24, 60, 64].

We have proved that it is impossible to completely forgo instrumentation in a HyTM even if only sequential TM-progress is required, and that any opaque HyTM implementation providing non-trivial progress either has to pay a *linear* number of metadata accesses, or will have to allow slow-path transactions to abort fast-path operations. The main motivation for our definition of metadata base objects (Definition 7.1) is given by experiments suggesting that the cost of concurrency detection is a significant bottleneck for many HyTM implementations [102]. To precisely characterize the costs incurred by hardware transactions, we made a distinction between the set of memory locations that store the data values of the t-objects and the locations that store the metadata information. To the best of our knowledge, all known HyTM proposals, such as *HybridNOrec* [35, 112], *PhTM* [99] and others [37, 88] avoid co-locating the data and metadata within a single base object.

HyTM algorithms. Circa 2005, several papers introduced HyTM implementations [12, 37, 88] that integrated HTMs with variants of *DSTM* [79]. These implementations provide nontrivial concurrency between hardware and software transactions (progressiveness), by imposing instrumentation on hardware transactions: every t-read operation incurs at least one extra access to a metadata base object. Our Theorem 7.9 shows that this overhead is unavoidable. Of note, write operations of these HyTMs are also instrumented, but our Algorithm 7.1 shows that it is not necessary.

Implementations like *PhTM* [99] and *HybridNOrec* [35] overcome the per-access instrumentation cost of [37, 88] by realizing that if one is prepared to sacrifice progress, hardware transactions need instrumentation only at the boundaries of transactions to detect pending software transactions. Inspired by

this observation, our HyTM implementation described in Algorithm 7.2 overcomes the linear per-read instrumentation cost by allowing hardware readers to abort due to a concurrent software writer, but maintains progressiveness for software transactions, unlike [35, 99, 102].

References [69, 112] provide detailed overviews on HyTM designs and implementations. The software component of the HyTM algorithms presented in this paper is inspired by progressive STM implementations [36, 39, 90] and is subject to the lower bounds for progressive STMs established in [23, 62, 64, 90].

8

Optimism for boosting concurrency

> The wickedness and the foolishness of no man can avail against the fond optimism of mankind.
>
> *James Branch Cabell*-The Silver Stallion

8.1 Overview

In previous chapters, we were concerned with the inherent complexities of implementing TM. In this chapter, we are concerned with using TM to derive concurrent implementations and raise a fundamental question about the ability of the TM abstraction to transform a sequential implementation to a concurrent one. Specifically, does the optimistic nature of TM give it an inherent advantage in exploiting concurrency that is lacking in pessimistic synchronization techniques like locking? To exploit concurrency, conventional lock-based synchronization pessimistically protects accesses to the shared memory before executing them. Speculative synchronization, achieved using TMs, optimistically executes memory operations with a risk of aborting them in the future. A programmer typically uses these synchronization techniques as "wrappers" to allow every process (or thread) to *locally* run its sequential code while ensuring that the resulting concurrent execution is *globally* correct.

Unfortunately, it is difficult for programmers to tell in advance which of the synchronization techniques will establish more concurrency in their resulting programs. In this chapter, we analyze the "amount of concurrency" one can obtain by turning a sequential program into a concurrent one. In particular, we compare the use of optimistic and pessimistic synchronization techniques, whose prime examples are TMs and locks respectively.

To fairly compare concurrency provided by implementations based on various techniques, one has (1) to define what it means for a concurrent program to be correct regardless of the type of synchronization it uses and (2) to define a metric of concurrency.

Correctness. We begin by defining a consistency criterion, namely *locally-serializable linearizability*. We say that a concurrent implementation of a given sequential data type is *locally serializable* if it ensures that the local execution of each operation is equivalent to *some* execution of its sequential implementation.

This condition is weaker than serializability since it does not require that there exists a *single* sequential execution that is consistent with all local executions. It is however sufficient to guarantee that optimistic executions do not observe an inconsistent transient state that could lead, for example, to a fatal error like division-by-zero.

Furthermore, the implementation should "make sense" globally, given the *sequential type* of the data structure we implement. The high-level history of every execution of a concurrent implementation must be *linearizable* [27, 83] with respect to this sequential type. The combination of local serializability and linearizability gives a correctness criterion that we call *LS-linearizability*, where LS stands for "locally serializable". We show that LS-linearizability is, as the original linearizability, compositional [81, 83]: a composition of LS-linearizable implementations is also LS-linearizable.

We apply the criterion of LS-linearizability to two broad classes of *pessimistic* and *optimistic* synchronization techniques. Pessimistic implementations capture what can be achieved using classic locks; in contrast, optimistic implementations proceed speculatively and fail to return a response to the process in the case of conflicts, *e.g.*, relying on transactional memory.

Measuring concurrency. We characterize the amount of concurrency provided by an LS-linearizable implementation as the set of schedules it accepts. To this end, we define a concurrency metric inspired by the analysis of parallelism in database concurrency control [74, 126]. More specifically, we assume an external scheduler that defines which processes execute which steps of the corresponding sequential program in a dynamic and unpredictable fashion. This allows us to define concurrency provided by an implementation as the set of *schedules* (interleavings of steps of concurrent sequential operations) it *accepts* (is able to effectively process). Then, the more schedules the implementation would accept, the more concurrent it would be.

We provide a framework to compare the concurrency one can get by choosing a particular synchronization technique for a specific data type. For the first time, we analytically capture the inherent concurrency provided by optimism-based and pessimism-based implementations in exploiting concurrency. We illustrate this using a popular sequential list-based set implementation [81], concurrent implementations of which are our running examples. More precisely, we show that there exist TM-based implementations that, for some workloads, allow for more concurrency than *any* pessimistic implementation, but we also show that there exist pessimistic implementations that, for other workloads, allow for more concurrency than *any* TM-based implementation.

Intuitively, an implementation based on transactions may abort an operation based on the way concurrent steps are scheduled, while a pessimistic implementation has to proceed eagerly without knowing about how future steps will be scheduled, sometimes over-conservatively rejecting a potentially acceptable schedule. By contrast, pessimistic implementations designed to exploit the semantics of the data type can supersede the "semantics-oblivious" TM-based implementations. More surprisingly, we demonstrate that combining the benefit of pessimistic implementations, namely their semantics awareness, and the benefit of TMs, namely their optimism, enables implementations that are strictly better-suited for exploiting concurrency than any of them individually. We describe a generic optimistic implementation of the list-based set that is *optimal* with respect to our concurrency metric: we show that, essentially, it accepts *all* correct concurrent schedules.

Our results suggest that "relaxed" TM models that are designed with the semantics of the high-level object in mind might be central to exploiting concurrency.

Roadmap of Chapter 8. In Section 8.2, we introduce the class of optimistic and pessimistic concurrent implementations we consider in this chapter. Section 8.3 introduces the definition of locally serializable linearizability and Section 8.4 is devoted to the concurrency analysis of optimistic and pessimistic synchronization techniques in the context of the list-based set. We wrap up with concluding remarks in Section 8.5.

8.2 Concurrent implementations

Objects and implementations. As with Chapter 2, we assume an asynchronous shared-memory system in which a set of $n > 1$ processes $p_1, \ldots, p_n$ communicate by applying *operations* on shared *objects*.

An object is an instance of an *abstract data type* which specifies a set of operations that provide the only means to manipulate the object. Recall that an *abstract data type* τ is a tuple $(\Phi, \Gamma, Q, q_0, \delta)$ where Φ is a set of operations, Γ is a set of responses, Q is a set of states, $q_0 \in Q$ is an initial state and $\delta \subseteq Q \times \Phi \times Q \times \Gamma$ is a transition relation that determines, for each state and each operation, the set of possible resulting states and produced responses. In this chapter, we consider only types that are *total*, *i.e.*, for every $q \in Q$, $\pi \in \Phi$, there exist $q' \in Q$ and $r \in \Gamma$ such that $(q, \pi, q', r) \in \delta$. We assume that every type $\tau = (\Phi, \Gamma, Q, q_0, \delta)$ is *computable*, i.e., there exists a Turing machine that, for each input (q, π), $q \in Q$, $\pi \in \Phi$, computes a pair (q', r) such that $(q, \pi, q', r) \in \delta$.

For any type τ, each high-level object O_τ of this type has a *sequential implementation*. For each operation $\pi \in \Phi$, *IS* specifies a deterministic procedure that performs *reads* and *writes* on a collection of objects $X_1, \ldots, X_m$ that encode a state of O_τ, and returns a response $r \in \Gamma$.

Sequential list-based set. As a running example, we consider the sorted linked-list based implementation of the type *set*, commonly referred to as the list-based set [81]. Recall that the set type exports operations $\mathsf{insert}(v)$, $\mathsf{remove}(v)$ and $\mathsf{contains}(v)$, with $v \in \mathbb{Z}$. Formally, the *set* type is defined by the tuple $(\Phi, \Gamma, Q, q_0, \delta)$ where:

Φ $= \{\mathsf{insert}(v), \mathsf{remove}(v), \mathsf{contains}(v)\}$; $v \in \mathbb{Z}$

Γ $= \{true, false\}$

Q is the set of all finite subsets of $\mathbb{Z}$; $q_0 = \emptyset$

δ is defined as follows:

(1): $(q, \mathsf{contains}(v), q, (v \in q))$

(2): $(q, \mathsf{insert}(v), q \cup \{v\}, (v \notin q))$

(3): $(q, \mathsf{remove}(v), q \setminus \{v\}, (v \in q))$

We consider a sequential implementation *LL* (Algorithm 8.1) of the set type using a sorted linked list where each element (or *object*) stores an integer value, *val*, and a pointer to its successor, *next*, so that elements are sorted in the ascending order of their value.

Every operation invoked with a parameter v traverses the list starting from the *head* up to the element storing value $v' \geq v$. If $v' = v$, then $\mathsf{contains}(v)$ returns true, $\mathsf{remove}(v)$ unlinks the corresponding element and returns true, and $\mathsf{insert}(v)$ returns false. Otherwise, $\mathsf{contains}(v)$ and $\mathsf{remove}(v)$ return false while $\mathsf{insert}(v)$ adds a new element with value v to the list and returns true. The list-based set is denoted by (LL, set).

Concurrent implementations. We tackle the problem of turning the sequential implementation *IS* of type τ into a *concurrent* one, shared by n processes. The implementation provides the processes with algorithms for the reads and writes on objects. We refer to the resulting implementation as a concurrent implementation of (IS, τ). As in Chapter 2, we assume an asynchronous shared-memory system in which the processes communicate by applying primitives on shared *base objects* [75]. We place no upper bounds on the number of versions an object may maintain or on the size of this object.

Throughout this chapter, the term *operation* refers to some high-level operation of the type, while read-write operations on objects are referred simply as *reads* and *writes*.

An implemented read or write may *abort* by returning a special response $\perp$. In this case, we say that the corresponding high-level operation is *aborted*. The $\perp$ event is treated both as the response event of the read or write operation and as the response of the corresponding high-level operation.

Algorithm 8.1 Sequential implementation *LL* (*sorted linked list*) of *set* type

```
1:  Shared variables:
2:    Initially head, tail,
3:      head.val = −∞, tail.val = +∞
4:      head.next = tail

5:  insert(v):
6:    prev ← head // copy the address
7:    curr ← read(prev.next) // fetch the next element
8:    while (tval ← read(curr.val)) < v do
9:      prev ← curr
10:     curr ← read(curr.next) // fetch from memory
11:   end while
12:   if tval ≠ v then // tval is stored locally
13:     X ← new-node(v, prev.next) // v and address of curr
14:     write(prev.next, X) // next points to the new element
15:   Return (tval ≠ v)
16: remove(v):
17:   prev ← head // copy the address
18:   curr ← read(prev.next) // fetch next field
19:   while (tval ← read(curr.val)) < v do // val local copy
20:     prev ← curr
21:     curr ← read(curr.next)
22:   end while
23:   if tval = v then
24:     tnext ← read(curr.next) // fetch the node after curr
25:     write(prev.next, tnext) // delete the node
26:   Return (tval = v)
27: contains(v):
28:   curr ← head
29:   curr ← read(prev.next)
30:   while (tval ← read(curr.val)) < v do
31:     curr ← read(curr.next)
32:   end while
33:   Return (tval = v)
```

Executions and histories. An *execution* of a concurrent implementation (of (IS, τ)) is a sequence of invocations and responses of high-level operations of type τ, invocations and responses of read and write operations, and primitives applied on base-objects. We assume that executions are *well-formed*: no process invokes a new read or write, or high-level operation before the previous read or write, or a high-level operation, resp., returns, or takes steps outside its read or write operation's interval.

Let $\alpha|p_i$ denote the subsequence of an execution α restricted to the events of process p_i. Executions α and α' are *equivalent* if for every process p_i, $\alpha|p_i = \alpha'|p_i$. An operation π *precedes* another operation π' in an execution α, denoted $\pi \rightarrow_\alpha \pi'$, if the response of π occurs before the invocation of π'. Two operations are *concurrent* if neither precedes the other. An execution is *sequential* if it has no concurrent operations. A sequential execution α is *legal* if for every object X, every read of X in α returns the latest written value of X. An operation is *complete* in α if the invocation event is followed by a *matching* (non-$\perp$) response or aborted; otherwise, it is *incomplete* in α. Execution α is *complete* if every operation is complete in α.

The *history exported by an execution* α is the subsequence of α reduced to the invocations and responses of operations, reads and writes, except for the reads and writes that return $\perp$.

High-level histories and linearizability. A *high-level history* $\tilde{H}$ of an execution α is the subsequence of α consisting of all invocations and responses of (high-level) operations.

Definition 8.1 (Linearizability)**.** *A complete high-level history* $\tilde{H}$ *is* linearizable *with respect to an object type* τ *if there exists a sequential high-level history* S *equivalent to* $\tilde{H}$ *such that (1)* $\rightarrow_{\tilde{H}} \subseteq \rightarrow_S$ *and (2)* S is consistent with the sequential specification of type τ.

Now a high-level history $\tilde{H}$ *is linearizable if it can be* completed *(by adding matching responses to a subset of incomplete operations in* $\tilde{H}$ *and removing the rest) to a linearizable high-level history [27, 83].*

Obedient implementations. We only consider implementations that satisfy the following condition: Let α be any complete sequential execution of a concurrent implementation I. Then in every execution of I of the form $\alpha \cdot \rho_1 \cdots \rho_k$ where each ρ_i $(i = 1, \ldots, k)$ is the complete execution of a read, every read returns the value written by the last write that does not belong to an aborted operation.

Intuitively, this assumption restricts our scope to "obedient" implementations of reads and writes, where no read value may depend on some future write. In particular, we filter out implementations in which the complete execution of a high-level operation is performed within the first read or write of its sequential implementation.

Pessimistic implementations. Informally, a concurrent implementation is *pessimistic* if the exported history contains every read-write event that appears in the execution. More precisely, no execution of a pessimistic implementation includes operations that returned $\perp$.

For example, a class of pessimistic implementations are those based on locks. A lock provides shared or exclusive access to an object X through synchronization primitives $\mathsf{lock}^S(X)$ (*shared mode*), $\mathsf{lock}(X)$ (*exclusive mode*), and $\mathsf{unlock}(X)$. When $\mathsf{lock}^S(X)$ (resp. $\mathsf{lock}(X)$) invoked by a process p_i returns, we say that p_i *holds a lock on* X *in shared (resp. exclusive) mode.* A process *releases* the object it holds by invoking $\mathsf{unlock}(X)$. If no process holds a shared or exclusive lock on X, then $\mathsf{lock}(X)$ eventually returns; if no process holds an exclusive lock on X, then $\mathsf{lock}^S(X)$ eventually returns; and if no process holds a lock on X forever, then every $\mathsf{lock}(X)$ or $\mathsf{lock}^S(X)$ eventually returns. Given a sequential implementation of a data type, a corresponding lock-based concurrent one is derived by inserting the synchronization primitives to provide read-write access to an object.

Optimistic implementations. In contrast with pessimistic ones, optimistic implementations may, under certain conditions, abort an operation: some read or write may return $\perp$, in which case the corresponding operation also returns $\perp$.

Popular classes of optimistic implementations are those based on "lazy synchronization" [71, 81] (with the ability of returning $\perp$ and re-invoking an operation) or transactional memory.

8.3 Locally serializable linearizability

We are now ready to define the correctness criterion that we impose on our concurrent implementations.

Let H be a history and let π be a high-level operation in H. Then $H|\pi$ denotes the subsequence of H consisting of the events of π, except for the last aborted read or write, if any. Let IS be a sequential implementation of an object of type τ and Σ_{IS}, the set of histories of IS.

Definition 8.2 (LS-linearizability)**.** *A history H is* locally serializable with respect to *IS if for every high-level operation π in H, there exists $S \in \Sigma_{IS}$ such that $H|\pi = S|\pi$. A history H is* LS-linearizable with respect to (IS, τ) *(we also write H is (IS, τ)-LSL) if: (1) H is locally serializable with respect to IS and (2) the corresponding high-level history $\tilde{H}$ is linearizable with respect to τ.*

Observe that local serializability stipulates that the execution is witnessed sequential by every operation. Two different operations (even when invoked by the same process) are not required to witness mutually consistent sequential executions.

A concurrent implementation I is *LS-linearizable with respect to* (IS, τ) (we also write I is (IS, τ)-LSL) if every history exported by I is (IS, τ)-LSL. Throughout this paper, when we refer to a concurrent implementation of (IS, τ), we assume that it is LS-linearizable with respect to (IS, τ).

LS-linearizability is compositional. Just as linearizability, LS-linearizability is *compositional* [81, 83]: a composition of LSL implementations is also LSL. We define the composition of two distinct object types τ_1 and τ_2 as a type $\tau_1 \times \tau_2 = (\Phi, \Gamma, Q, q_0, \delta)$ as follows: $\Phi = \Phi_1 \cup \Phi_2$, $\Gamma = \Gamma_1 \cup \Gamma_2$,[1] $Q = Q_1 \times Q_2$,

[1] Here we treat each τ_i as a distinct type by adding index i to all elements of Φ_i, Γ_i, and Q_i.

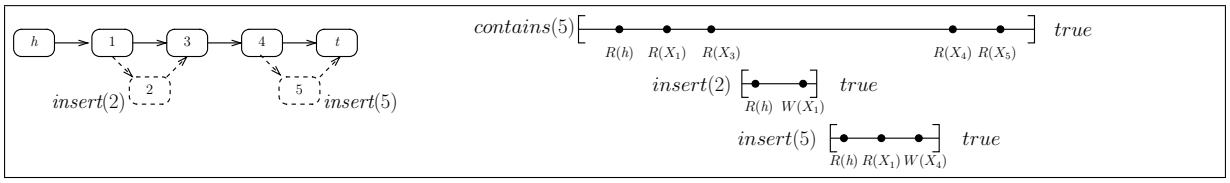

Figure 8.1: A concurrency scenario for a list-based set, initially $\{1, 3, 4\}$, where value i is stored at node X_i: insert(2) and insert(5) can proceed concurrently with contains(5), the history is LS-linearizable but not serializable. (We only depict important read-write events here.)

$q_0 = (q_{01}, q_{02})$, and $\delta \subseteq Q \times \Phi \times Q \times \Gamma$ is such that $((q_1, q_2), \pi, (q_1' q_2'), r) \in \delta$ if and only if for $i \in \{1, 2\}$, if $\pi \in \Phi_i$ then $(q_i, \pi, q_i', r) \in \delta_i \wedge q_{3-i} = q_{3-i}'$.

Every sequential implementation IS of an object $O_1 \times O_2$ of a composed type $\tau_1 \times \tau_2$ naturally induces two sequential implementations I_{S1} and I_{S2} of objects O_1 and O_2, respectively. Now a correctness criterion Ψ is *compositional* if for every history H on an object composition $O_1 \times O_2$, if Ψ holds for $H|O_i$ with respect to I_{Si}, for $i \in \{1, 2\}$, then Ψ holds for H with respect to $IS = I_{S1} \times I_{S2}$. Here, $H|O_i$ denotes the subsequence of H consisting of events on O_i.

Theorem 8.1. *LS-linearizability is compositional.*

Proof. Let H, a history on $O_1 \times O_2$, be LS-linearizable with respect to IS. Let each $H|O_i$, $i \in \{1, 2\}$, be LS-linearizable with respect to I_{Si}. Without loss of generality, we assume that H is complete (if H is incomplete, we consider any completion of it containing LS-linearizable completions of $H|O_1$ and $H|O_1$).

Let $\tilde{H}$ be a completion of the high-level history corresponding to H such that $\tilde{H}|O_1$ and $\tilde{H}|O_2$ are linearizable with respect to τ_1 and τ_2, respectively. Since linearizability is compositional [81, 83], $\tilde{H}$ is linearizable with respect to $\tau_1 \times \tau_2$.

Now let, for each operation π, S_π^1 and S_π^2 be any two sequential histories of I_{S1} and I_{S2} such that $H|\pi|O_j = S_\pi^j|\pi$, for $j \in \{1, 2\}$ (since $H|O_1$ and $H|O_2$ are LS-linearizable such histories exist). We construct a sequential history S_π by interleaving events of S_π^1 and S_π^2 so that $S_\pi|O_j = S_\pi^j$, $j \in \{1, 2\}$. Since each S_π^j acts on a distinct component O_j of $O_1 \times O_2$, every such S_π is a sequential history of IS. We pick one S_π that respects the local history $H|\pi$, which is possible, since $H|\pi$ is consistent with both $S_1|\pi$ and $S_2|\pi$.

Thus, for each π, we obtain a history of IS that agrees with $H|\pi$. Moreover, the high-level history of H is linearizable. Thus, H is LS-linearizable with respect to IS. □

LS-linearizability versus other criteria. LS-linearizability is a two-level consistency criterion which makes it suitable to compare concurrent implementations of a sequential data structure, regardless of synchronization techniques they use. It is quite distinct from related criteria designed for database and software transactions, such as serializability [109, 125] and multilevel serializability [124, 125].

For example, serializability [109] prevents sequences of reads and writes from conflicting in a cyclic way, establishing a global order of transactions. Reasoning only at the level of reads and writes may be overly conservative: higher-level operations may commute even if their reads and writes conflict [123]. Consider an execution of a concurrent *list-based set* depicted in Figure 8.1. We assume here that the set initial state is $\{1, 3, 4\}$. Operation contains(5) is concurrent, first with operation insert(2) and then with operation insert(5). The history is not serializable: insert(5) sees the effect of insert(2) because $R(X_1)$ by insert(5) returns the value of X_1 that is updated by insert(2) and thus should be serialized after it. But contains(5) misses element 2 in the linked list, but must see the effect of insert(5) to perform the read of X_5, i.e., the element created by insert(5). However, this history is LSL since each of the three local histories is consistent with some sequential history of LL.

Multilevel serializability [124, 125] was proposed to reason in terms of multiple semantic levels in the same execution. LS-linearizability, being defined for two levels only, does not require a global serialization of low-level operations as 2-level serializability does. LS-linearizability simply requires each

process to observe a local serialization, which can be different from one process to another. Also, to make it more suitable for concurrency analysis of a concrete data structure, instead of semantic-based commutativity [123], we use the sequential specification of the high-level behavior of the object [83].

Linearizability [27, 83] only accounts for high-level behavior of a data structure, so it does not imply LS-linearizability. For example, Herlihy's universal construction [75] provides a linearizable implementation for any given object type, but does not guarantee that each execution locally appears sequential with respect to any sequential implementation of the type. Local serializability, by itself, does not require any synchronization between processes and can be trivially implemented without communication among the processes. Therefore, the two parts of LS-linearizability indeed complement each other.

8.4 Pessimistic vs. optimistic synchronization

In this section, we compare the relative abilities of optimistic and pessimistic synchronization techniques to exploit concurrency in the context of the list-based set.

To characterize the ability of a concurrent implementation to process arbitrary interleavings of sequential code, we introduce the notion of a *schedule*. Intuitively, a schedule describes the order in which complete high-level operations, and *sequential* reads and writes are invoked by the user. More precisely, a schedule is an equivalence class of complete histories that agree on the *order* of invocation and response events of reads, writes and high-level operations, but not necessarily on read *values* or high-level responses. Thus, a schedule can be treated as a history, where responses of reads and operations are not specified.

We say that an implementation I *accepts* a schedule σ if it exports a history H such that $complete(H)$ exhibits the order of σ, where $complete(H)$ is the subsequence of H that consists of the events of the complete operations that returned a matching response. We then say that the execution (or history) *exports* σ. A schedule σ is (IS, τ)-LSL if there exists an (IS, τ)-LSL history that exports σ.

A *synchronization technique* is a set of concurrent implementations. We define a specific optimistic synchronization technique and then a specific pessimistic one.

The class $\mathcal{SM}$. Formally, $\mathcal{SM}$ denotes the set of optimistic, safe-strict serializable LSL implementations.

Let α denote the execution of a concurrent implementation and $ops(\alpha)$, the set of operations each of which performs at least one event in α. Let α^k denote the prefix of α up to the last event of operation π_k. Let $Cseq(\alpha)$ denote the set of subsequences of α that consist of all the events of operations that are complete in α. We say that α is *strictly serializable* if there exists a legal sequential execution α' equivalent to a sequence in $\sigma \in Cseq(\alpha)$ such that $\rightarrow_\sigma \subseteq \rightarrow_{\alpha'}$.

We focus on optimistic implementations that are strictly serializable and, in addition, guarantee that every operation (even aborted or incomplete) observes correct (serial) behavior. More precisely, an execution α is *safe-strict serializable* if (1) α is strictly serializable, and (2) for each operation π_k, there exists a legal sequential execution $\alpha' = \pi_0 \cdots \pi_i \cdot \pi_k$ and $\sigma \in Cseq(\alpha^k)$ such that $\{\pi_0, \cdots, \pi_i\} \subseteq ops(\sigma)$ and $\forall \pi_m \in ops(\alpha') : \alpha'|m = \alpha^k|m$.

Similar to other relaxations of opacity [64] like *TMS1* [43] and *VWC* [85], safe-strict serializable implementations ($\mathcal{SM}$) require that every transaction (even aborted and incomplete) observes "correct" serial behavior. Safe-strict serializability captures nicely both local serializability and linearizability. If we transform a sequential implementation IS of a type τ into a *safe-strict serializable* concurrent one, we obtain an LSL implementation of (IS, τ). Thus, the following lemma is immediate.

Lemma 8.2. *Let I be a safe-strict serializable implementation of (IS, τ). Then, I is LS-linearizable with respect to (IS, τ).*

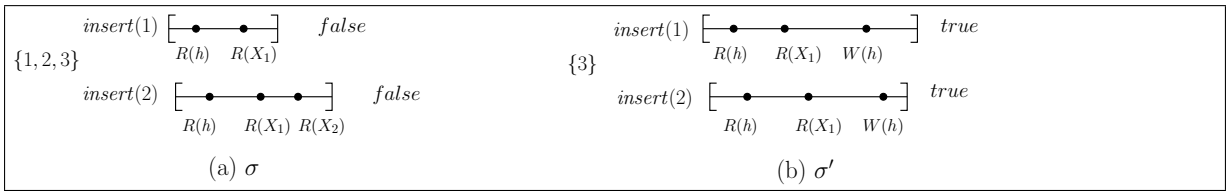

Figure 8.2: (a) a history exporting schedule σ, with initial state $\{1, 2, 3\}$, accepted by $I^{LP} \in \mathcal{SM}$; (b) a history exporting a problematic schedule σ', with initial state $\{3\}$, which should be accepted by any $I \in \mathcal{P}$ if it accepts σ

Indeed, by running each operation of *IS* within a transaction of a safe-strict serializable TM, we make sure that completed operations witness the same execution of *IS*, and every operation that returned $\perp$ is consistent with some execution of *IS* based on previously completed operations.

The class $\mathcal{P}$. This denotes the set of *deadlock-free* pessimistic LSL implementations: assuming that every process takes enough steps, at least one of the concurrent operations return a matching response [82]. Note that $\mathcal{P}$ includes implementations that are not necessarily safe-strict serializable.

8.4.1 Concurrency analysis

We now provide a concurrency analysis of synchronization techniques $\mathcal{SM}$ and $\mathcal{P}$ in the context of the list-based set.

A pessimistic implementation $I^H \in \mathcal{P}$ of (LL, set). We describe a pessimistic implementation of (LL, set), $I^H \in \mathcal{P}$, that accepts non-serializable schedules: each read operation performed by contains acquires the *shared lock* on the object, reads the *next* field of the element before releasing the shared lock on the predecessor element in a *hand-over-hand* manner [29]. Update operations (insert and remove) acquire the *exclusive lock* on the *head* during read(*head*) and release it at the end. Every other read operation performed by update operations simply reads the element *next* field to traverse the list. The write operation performed by an insert or a remove acquires the exclusive lock, writes the value to the element and releases the lock. There is no real concurrency between any two update operations since the process holds the exclusive lock on the *head* throughout the operation execution. Thus:

Lemma 8.3. *I^H is deadlock-free and LSL implementation of (LL, set).*

On the one hand, the schedule of (LL, set) depicted in Figure 8.1, which we denote by σ_0, is not serializable and must be rejected by any implementation in $\mathcal{SM}$. However, there exists an execution of I^H that exports σ_0 since there is no read-write conflict on any two consecutive elements accessed.

On the other hand, consider the schedule σ of (LL, set) in Figure 8.2(a). Clearly, σ is serializable and is accepted by implementations based on most progressive TMs since there is no read-write conflict. For example, let I^{LP} denote an implementation of (IS, τ) based on the progressive opaque TM implementation LP in Algorithm 4.1 (Chapter 4). Then, $I^{LP} \in \mathcal{SM}$ and the schedule σ is accepted by I^{LP}. However, we prove that σ is not accepted by any implementation in $\mathcal{P}$. Our proof technique is interesting in its own right: we show that if there exists any implementation in $\mathcal{P}$ that accepts σ, it must also accept the schedule σ' depicted in Figure 8.2(b). In σ', insert(2) overwrites the write on *head* performed by insert(1) resulting in a lost update. By deadlock-freedom, there exists an extension of σ' in which a contains(1) returns *false*; but this is not a linearizable schedule.

Theorem 8.4. *There exists a schedule σ_0 of (LL, set) that is accepted by an implementation in $\mathcal{PL}$, but not accepted by* any *implementation $I \in \mathcal{SM}$.*

Proof. Let σ_0 be the schedule of (LL, set) depicted in Figure 8.1. Suppose by contradiction that $\sigma_0 \in \mathcal{S}(I)$, where I is an implementation of (LL, set) based on any safe-strict serializable TM. Thus, there exists an execution α of I that exports σ_0. Now consider two cases: (1) Suppose that the read of X_4 by

contains(5) returns the value of X_4 that is updated by insert(5). Since insert(2) $\rightarrow_\alpha$ insert(5), insert(2) must precede insert(5) in any sequential execution α' equivalent to α. Also, since contains(5) reads X_1 prior to its update by insert(2), contains(5) must precede insert(2) in α'. But then the read of X_4 is not legal in α'—a contradiction since α must be serializable. (2) Suppose that contains(5) reads the initial value of X_4, *i.e.*, its value prior to the write to X_4 by insert(5), where $X_4.next$ points to the *tail* of the list (according to our sequential implementation LL). But then, according to LL, contains(5) cannot access X_5 in σ_0—a contradiction.

Consider the pessimistic implementation $I^H \in \mathcal{P}$: since the contains operation traverses the list using shared hand-over-hand locking, the process p_i executing contains(5) can release the lock on element X_1 prior to the acquisition of the exclusive lock on X_1 by insert(2). Similarly, p_i can acquire the shared lock on X_4 immediately after the release of the exclusive lock on X_4 by the process executing insert(5) while still holding the shared lock on element X_3. Thus, there exists an execution of I^H that exports σ_0. □

Theorem 8.5. *There exists a schedule σ of (LL, set) that is accepted by an implementation in $\mathcal{SM}$, but not accepted by* any *implementation in $\mathcal{P}$.*

Proof. We show first that the schedule σ of (LL, set) depicted in Figure 8.2(a) is not accepted by any implementation in $\mathcal{P}$. Suppose the contrary and let σ be exported by an execution α. Here α starts with three sequential insert operations with parameters 1, 2, and 3. The resulting "state" of the set is $\{1, 2, 3\}$, where value $i \in \{1, 2, 3\}$ is stored in object X_i.

Suppose, by contradiction, that some $I \in \mathcal{P}$ accepts σ. We show that I then accepts the schedule σ' depicted in Figure 8.2(b), which starts with a sequential execution of insert(3) storing value 3 in object X_1.

Let α' be any history of I that exports σ'. Recall that we only consider obedient implementations: in α': the read of *head* by insert(2) in σ' refers to X_1 (the next element to be read by insert(2)). In α, element X_1 stores value 1, *i.e.*, insert(1) can safely return false, while in σ', X_1 stores value 3, i.e., the next step of insert(1) must be a write to *head*. Thus, no process can distinguish α and α' before the read operations on X_1 return. Let α'' be the prefix of α' ending with $R(X_1)$ executed by insert(2). Since I is deadlock-free, we have an extension of α'' in which both insert(1) and insert(2) terminate; we show that this extension violates linearizability. Since I is locally-serializable, to respect our sequential implementation of (LL, set), both operations should complete the write to *head* before returning. Let π_1 = insert(1) be the first operation to write to *head* in this extended execution. Let π_2 = insert(2) be the other insert operation. It is clear that π_1 returns *true* even though π_2 overwrites the update of π_1 on *head* and also returns *true*. Recall that implementations in $\mathcal{P}$ are deadlock-free. Thus, we can further extend the execution with a complete contains(1) that will return *false* (the element inserted to the list by π_1 is lost)—a contradiction since I is linearizable with respect to *set*. Thus, $\sigma \notin \mathcal{S}(I)$ for any $I \in \mathcal{P}$.

On the other hand, the schedule σ is accepted by $I^{LP} \in \mathcal{SM}$, since there is no conflict between the two concurrent update operations. □

8.4.2 Concurrency optimality

We now combine the benefits of semantics awareness of implementations in $\mathcal{P}$ and the optimism of $\mathcal{SM}$ to derive a generic optimistic implementation of the list-based set that supersedes every implementation in classes $\mathcal{P}$ and $\mathcal{SM}$ in terms of concurrency. Our implementation, denoted I^{RM} provides processes with algorithms for implementing read and write operations on the elements of the list for each operation of the list-based set (Algorithm 8.2).

Every object (or element) X_ℓ is specified by the following shared variables: $t\text{-}var[\ell]$ stores the *value* $v \in V$ of X_ℓ, $r[\ell]$ stores a boolean indicating if X_ℓ is *marked for deletion*, $L[\ell]$ stores a tuple of the *version number* of X_ℓ and a *locked* flag; the latter indicates whether a concurrent process is performing a write to X_ℓ.

Algorithm 8.2 Code for process p_k implementing reads and writes in implementation I^{RM}

```
1:  Shared variables:
2:    for each object X_ℓ:
3:          t-var[ℓ], initially 0
4:          r[ℓ], initially false
5:          L[ℓ] ∈ 𝒩 × {true, false} supports read
6:               write, cas operations, initially ⟨0, false⟩

7:  Local variables of process p_k:
8:    rbuf_k[i] ⊂ X × 𝒩; i = {1, 2} cyclic buffer of size 2,
9:       initially ∅

10: read_k(X_ℓ) executed by insert, remove, contains:
11:   ⟨ver_1, ∗⟩ ← L[ℓ].read() // get versioned lock
12:   val ← t-var[ℓ].read() // get value
13:   r ← r[ℓ].read()
14:   ⟨ver_2, ∗⟩ ← L[ℓ].read() // reget versioned lock
15:   if (ver_1 ≠ ver_2) ∨ r then
16:      Return ⊥
17:   rbuf_k.add(⟨X_ℓ, ver_1⟩) // override penultimate entry
18:   Return val

19: write_k(X_ℓ, v) executed by remove:
20:   let oldver_ℓ be such that ⟨X_ℓ, oldver_ℓ⟩ ∈ rbuf_k
21:   ver ← oldver_ℓ
22:   if ¬L[ℓ].cas(⟨ver, false⟩, ⟨ver, true⟩) then
23:      Return ⊥ // grab lock or abort
24:   let X_ℓ' ≠ X_ℓ be such that {X_ℓ', ver_ℓ'} ∈ rbuf_k
25:   if ¬L[ℓ'].cas(⟨ver_ℓ', false⟩, ⟨ver_ℓ', true⟩) then
26:      Return ⊥ // grab lock or abort
27:   r[ℓ'].write(true) // mark element for deletion
28:   t-var[ℓ].write(v) // update memory
29:   L[ℓ].write(⟨ver + 1, false⟩)// release locks
30:   L[ℓ'].write(⟨ver_ℓ' + 1, false⟩)
31:   Return ok

32: write_k(X_ℓ, v) executed by insert:
33:   let oldver_ℓ be such that ⟨X_ℓ, oldver_ℓ⟩ ∈ rbuf_k
34:   ver ← oldver_ℓ
35:   if ¬L[ℓ].cas(⟨ver, false⟩, ⟨ver, true⟩) then
36:      Return ⊥ // grab lock or abort
37:   t-var[ℓ].write(v) // update memory
38:   L[ℓ].write(⟨ver + 1, false⟩)// release locks
39:   Return ok
```

Any operation with input parameter v traverses the list starting from the *head* element up to the element storing value $v' \geq v$ without writing to shared memory. If a read operation on an element conflicts with a write operation to the same element or if the element is marked for deletion, the operation terminates by returning $\perp$. While traversing the list, the process maintains the last two read elements and their version numbers in the local rotating buffer *rbuf*. If none of the read operations performed by contains(v) return $\perp$ and if $v' = v$, then contains(v) returns true; otherwise it returns false. Thus, the contains does not write to shared memory.

To perform write operation to an element as part of an update operation (insert and remove), the process first retrieves the version of the object that belongs to its rotating buffer. It returns $\perp$ if the version has been changed since the previous read of the element or if a concurrent process is executing a write to the same element. Note that, technically, $\perp$ is returned only if $prev.next \not\rightarrow curr$. If $prev.next \rightarrow curr$, then we attempt to lock the element with the current version and return $\perp$ if there is a concurrent process executing a write to the same element. But we avoid expanding on this step in our algorithm pseudocode. The write operation performed by the remove operation, additionally checks if the element to be removed from the list is locked by another process; if not, it sets a flag on the element to mark it for deletion. If none of the read or write operations performed during the insert(v) or remove(v) returned $\perp$,

appropriate matching responses are returned as prescribed by the sequential implementation LL. Any update operation of I^{RM} uses at most two expensive synchronization patterns [17].

Proof of LS-linearizability. Let α be an execution of I^{RM} and $<_\alpha$ denote the total-order on events in α. For simplicity, we assume that α starts with an artificial sequential execution of an insert operation π_0 that inserts *tail* and sets $head.next = tail$. Let H be the history exported by α, where all reads and writes are sequential. We construct H by associating a linearization point ℓ_{op} with each non-aborted read or write operation *op* performed in α as follows:

- if *op* is a read, then performed by process p_k, ℓ_{op} is the base-object read in line 12;
- if *op* is a write within an insert operation, ℓ_{op} is the base-object cas in line 22;
- if *op* is a write within a remove operation, ℓ_{op} is the base-object cas in line 35.

We say that a read of an element X within an operation π is *valid* in H (we also say that X is *valid*) if there does not exist any remove operation π_1 that *deallocates* X (removes X from the list) such that $\ell_{\pi_1.\mathsf{write}(X)} <_\alpha \ell_{\pi.\mathsf{read}(X)}$.

Lemma 8.6. *Let π be any operation performing* read(X) *followed by* read(Y) *in H. Then (1) there exists an* insert *operation that sets $X.next = Y$ prior to π.*read(X)*, and (2) π.*read(X) *and π.*read(Y) *are* valid *in H.*

Proof. Let π be any operation in I^{RM} that performs read(X) followed by read(Y). If X and Y are *head* and *tail* respectively, $head.next = tail$ (by assumption). Since no remove operation deallocates the *head* or *tail*, the read of X and Y are *valid* in H.

Now, let X be the *head* element and suppose that π performs read(X) followed by read(Y); $Y \neq tail$ in H. Clearly, if π performs a read(Y), there exists an operation $\pi' =$ insert that has previously set $head.next = Y$. More specifically, π.read(X) performs the action in line 12 after the write to shared memory by π' in line 37. By the assignment of linearization points to tx-operations, $\ell_{\pi'} <_\alpha \ell_{\pi.\mathsf{read}(X)}$. Thus, there exists an insert operation that sets $X.next = Y$ prior to π.read(X) in H.

For the second claim, we need to prove that the read(Y) by π is *valid* in H. Suppose by contradiction that Y has been deallocated by some $\pi'' =$ remove operation prior to read(Y) by π. By the rules for linearization of read and write operations, the action in line 28 precedes the action in line 12. However, π proceeds to perform the check in line 15 and returns $\perp$ since the flag corresponding to the element Y is previously set by π''. Thus, H does not contain π.read(Y)—contradiction.

Inductively, by the above arguments, every non-*head* read by π is performed on an element previously created by an insert operation and is valid in H. □

Lemma 8.7. *H is locally serializable with respect to LL.*

Proof. By Lemma 8.6, every element X read within an operation π is previously created by an insert operation and is valid in H. Moreover, if the read operation on X returns v', then $X.next$ stores a pointer to another valid element that stores an integer value $v'' > v'$. Note that the series of reads performed by π terminates as soon as an element storing value v or higher is found. Thus, π performs at most $O(|v - v_0|)$ reads, where v_0 is the value of the second element read by π. Now we construct S^π as a sequence of insert operations, that insert values read by π, one by one, followed by π. By construction, $S^\pi \in \Sigma_{LL}$. □

It is sufficient for us to prove that every finite high-level history H of I^{RM} is linearizable. First, we obtain a completion $\tilde{H}$ of H as follows. The invocation of an incomplete *contains* operation is discarded. The invocation of an incomplete $\pi =$ insert $\vee$ remove operation that has not returned successfully from the write operation is discarded; otherwise, it is completed with response *true*.

We obtain a sequential high-level history $\tilde{S}$ equivalent to $\tilde{H}$ by associating a linearization point ℓ_π with each operation π as follows. For each $\pi =$ insert $\vee$ remove that returns *true* in $\tilde{H}$, ℓ_π is associated with the first write performed by π in H; otherwise ℓ_π is associated with the last read performed by π in H.

For $\pi = \mathsf{contains}$ that returns *true*, ℓ_π is associated with the last read performed in I^{RM}; otherwise ℓ_π is associated with the read of *head*. Since linearization points are chosen within the intervals of operations of I^{RM}, for any two operations π_i and π_j in $\tilde{H}$, if $\pi_i \rightarrow_{\tilde{H}} \pi_j$, then $\pi_i \rightarrow_{\tilde{S}} \pi_j$.

Lemma 8.8. *$\tilde{S}$ is consistent with the sequential specification of type set.*

Proof. Let $\tilde{S}^k$ be the prefix of $\tilde{S}$ consisting of the first k complete operations. We associate each $\tilde{S}^k$ with a set q^k of objects that were successfully inserted and not subsequently successfully removed in $\tilde{S}^k$. We show by induction on k that the sequence of state transitions in $\tilde{S}^k$ is consistent with operations' responses in $\tilde{S}^k$ with respect to the *set* type.

The base case $k = 1$ is trivial: the *tail* element containing $+\infty$ is successfully inserted. Suppose that $\tilde{S}^k$ is consistent with the *set* type and let π_1 with argument $v \in \mathbb{Z}$ and response r_{π_1} be the last operation of $\tilde{S}^{k+1}$. We want to show that $(q^k, \pi_1, q^{k+1}, r_{\pi_1})$ is consistent with the *set* type.

(1) If $\pi_1 = \mathsf{insert}(v)$ returns *true* in $\tilde{S}^{k+1}$, there does not exist any other $\pi_2 = \mathsf{insert}(v)$ that returns *true* in $\tilde{S}^{k+1}$ such that there does not exist any $\mathsf{remove}(v)$ that returns *true*; $\pi_2 \rightarrow_{\tilde{S}^{k+1}} \mathsf{remove}(v) \rightarrow_{\tilde{S}^{k+1}} \pi_1$. Suppose by contradiction that such a π_1 and π_2 exist. Every successful $\mathsf{insert}(v)$ operation performs its penultimate read on an element X that stores a value $v' < v$ and the last read is performed on an element that stores a value $v'' > v$. Clearly, π_1 also performs a write on X. By construction of $\tilde{S}$, π_1 is linearized at the release of the *cas* lock on element X. Observe that π_2 must also perform a write to the element X (otherwise one of π_1 or π_2 would return *false*). By assumption, the write to X in shared-memory by π_2 (line 37) precedes the corresponding write to X in shared-memory by π_2. If $\ell_{\pi_2} <_\alpha \ell_{\pi_1.\mathsf{read}(X)}$, then π_1 cannot return *true*—a contradiction. Otherwise, if $\ell_{\pi_1.\mathsf{read}(X)} <_\alpha \ell_{\pi_2}$, then π_1 reaches line 22 and return $\perp$. This is because either π_1 attempts to acquire the *cas* lock on X while it is still held by π_2 or the value of X contained in the *rbuf* of the process executing π_1 has changed—a contradiction.

If $\pi_1 = \mathsf{insert}(v)$ returns *false* in $\tilde{S}^{k+1}$, there exists a $\pi_2 = \mathsf{insert}(v)$ that returns *true* in $\tilde{S}^{k+1}$ such that there does not exist any $\pi_3 = \mathsf{remove}(v)$ that returns *true*; $\pi_2 \rightarrow_{\tilde{S}^{k+1}} \pi_3 \rightarrow_{\tilde{S}^{k+1}} \pi_1$. Suppose that such a π_2 does not exist. Thus, π_1 must perform its last read on an element that stores value $v'' > v$, perform the action in Line 37 and return *true*—a contradiction.

It is easy to verify that the conjunction of the above two claims prove that $\forall q \in Q$; $\forall v \in \mathbb{Z}$, $\tilde{S}^{k+1}$ satisfies $(q, \mathsf{insert}(v), q \cup \{v\}, (v \notin q))$.

(2) If $\pi_1 = \mathsf{remove}(v)$, similar arguments as applied to $\mathsf{insert}(v)$ prove that $\forall q \in Q$; $\forall v \in \mathbb{Z}$, $\tilde{S}^{k+1}$ satisfies $(q, \mathsf{remove}(v), q \setminus \{v\}, (v \in q))$.

(3) If $\pi_1 = \mathsf{contains}(v)$ returns *true* in $\tilde{S}^{k+1}$, there exists $\pi_2 = \mathsf{insert}(v)$ that returns *true* in $\tilde{S}^{k+1}$ such that there does not exist any $\mathsf{remove}(v)$ that returns *true* in $\tilde{S}^{k+1}$ such that $\pi_2 \rightarrow_{\tilde{S}^{k+1}} \mathsf{remove}(v) \rightarrow_{\tilde{S}^{k+1}} \pi_1$. The proof of this claim immediately follows from Lemma 8.6.

Now, if $\pi_1 = \mathsf{contains}(v)$ returns *false* in $\tilde{S}^{k+1}$, there does not exist an $\pi_2 = \mathsf{insert}(v)$ that returns *true* such that there does not exist any $\mathsf{remove}(v)$ that returns *true*; $\pi_2 \rightarrow_{\tilde{S}^{k+1}} \mathsf{remove}(v) \rightarrow_{\tilde{S}^{k+1}} \mathsf{contains}(v)$. Suppose by contradiction that such a π_1 and π_2 exist. Thus, the action in line 37 by the $\mathsf{insert}(v)$ operation that updates some element, say X precedes the action in line 12 by $\mathsf{contains}(v)$ that is associated with its first read (the *head*). We claim that $\mathsf{contains}(v)$ must read the element X' newly created by $\mathsf{insert}(v)$ and return *true*—a contradiction to the initial assumption that it returns *false*. The only case when this can happen is if there exists a remove operation that forces X' to be unreachable from *head* i.e. concurrent to the write to X by insert, there exists a remove that sets $X''.next$ to $X.next$ after the action in line 35 by insert. But this is not possible since the *cas* on X performed by the remove would return *false*.

Thus, inductively, the sequence of state transitions in $\tilde{S}$ satisfies the sequential specification of the *set* type. □

Lemmas 8.7 and 8.8 imply:

Theorem 8.9. *I^{RM} is LS-linearizable with respect to (LL, set).*

Proof of concurrency optimality. Now we show that I^{RM} supersedes, in terms of concurrency, *any* implementation in classes $\mathcal{P}$ or $\mathcal{SM}$. The proof is based on a more general optimality result, interesting in its own right: any finite schedule rejected by I^{RM} is not *observably LS-linearizable* (or simply *observable*). We show that any finite schedule rejected by our algorithm is not *observably correct*.

A correct schedule σ is observably correct if by completing update operations in σ and extending, for any $v \in \mathbb{Z}$, the resulting schedule with a complete sequential execution contains(v), applied to the resulting contents of the list, we obtain a correct schedule. Here the contents of the list after a given correct schedule is determined based on the order of its write operations. For each element, we define the resulting state of its *next* field based on the last write in the schedule. Since in a correct schedule, each new element is first created and then linked to the list, we can reconstruct the *state of the list* by iteratively traversing it, starting from *head*.

Intuitively, a schedule is observably correct if it incurs no "lost updates". Consider, for example a schedule (cf. Figure 8.2(b)) in which two operations, insert(1) and insert(2) are applied to the list with state {3}. The resulting schedule is trivially correct (both operations return *true* so the schedule can some from a complete linearizable history). However, in the schedule, one of the operations, say insert(1), overwrites the effect of the other one. Thus, if we extend the schedule with a complete execution of contains(2), the only possible response it may give is *false* which obviously does not produce a linearizable high-level history.

Theorem 8.10 (Optimality)**.** *I^{RM} accepts all schedules that are observable with respect to (LL, set).*

Proof. We prove that any schedule rejected by I^{RM} is not observable. We go through the cases when a read or write returns $\bot$ (implying the operation fails to return a matching response) and thus the current schedule is rejected: (1) read(X_ℓ) returns $\bot$ in line 16 when $r[\ell] = true$ or when $ver_1 \neq ver_2$, (2) write(X_ℓ) performed by remove(v) either returns $\bot$ in line 22 when the cas operation on $L[\ell]$ returns *false* or returns $\bot$ in line 25 when the cas operation on the element that stores v returns *false*, and (3) write(X_ℓ) performed by insert returns $\bot$ in line 35 when the cas operation on $L[\ell]$ returns *false*.

Consider the subcase (1a), $r[\ell]$ is set *true* by a preceding or concurrent write(X_ℓ) (line 27). The high-level operation performing this write is a remove that marks the corresponding list element as removed. Since no removed element can be read in a sequential execution of LL, the corresponding history is not locally serializable. Alternatively, in subcase (1b), the version of X_ℓ read previously in line 11 has changed. Thus, an update operation has concurrently performed a write to X_ℓ. However, there exist executions that export such schedules.

In case (2), the write performed by a remove operation returns $\bot$. In subcase (2a), X_ℓ is currently locked. Thus, a concurrent high-level operation has previously locked X_ℓ (by successfully performing $L[\ell]$.cas() in line 22) and has not yet released the lock (by writing $\langle ver', \mathsf{false}\rangle$ to $L[\ell]$ in line 29). In subcase (2b), the current version of X_ℓ (stored in $L[\ell]$) differs from the version of X_ℓ witnessed by a preceding read. Thus, a concurrent high-level operation completed a write to X_ℓ *after* the current high-level operation π performed a read of X_ℓ. In both (2a) and (2b), a concurrent high-level updating operation π' (remove or insert) has written or is about to perform a write to X_ℓ. In subcase (2c), the cas on the element $X_{\ell'}$ (element that stores the value v) executed by remove(v) returns *false* (line 25). Recall that by the sequential implementation LL, remove(v) performs a read of $X_{\ell'}$ prior to the write(X_ℓ), where $X_\ell.next$ refers to $X_{\ell'}$. If the *cas* on $X_{\ell'}$ fails, there exists a process that concurrently performed a write to $X_{\ell'}$, but after the read of $X_{\ell'}$ by remove(v). In all cases, we observe that if we did not abort the write to X_ℓ, then the schedule extended by a complete execution of contains is not LSL.

In case (3), the write performed by an insert operation returns $\bot$. Similar arguments to case (2) prove that any schedule rejected is not observable LSL. □

Theorem 8.10 implies that the schedules exported by the histories in Figures 8.1 and 8.2(a) and that are not accepted by any $I' \in \mathcal{SM}$ and any $I \in \mathcal{P}$, respectively, are indeed accepted by I^{RM}. But it is easy

to see that implementations in $\mathcal{SM}$ and $\mathcal{P}$ can only accept observable schedules. As a result, I^{RM} can be shown to strictly supersede any pessimistic or TM-based implementation of the list-based set.

Corollary 8.11. *I^{RM} accepts every schedule accepted by any implementation in $\mathcal{P}$ and $\mathcal{SM}$. Moreover, I^{RM} accepts schedules σ and σ' that are rejected by any implementation in $\mathcal{P}$ and $\mathcal{SM}$, respectively.*

8.5 Related work and Discussion

Measuring concurrency. Sets of accepted schedules are commonly used as a metric of concurrency provided by a shared memory implementation. Gramoli et al. [53] defined a concurrency metric, the *input acceptance*, as the ratio of committed transactions over aborted transactions when TM executes the given schedule. Unlike our metric, input acceptance does not apply to lock-based programs.

For static database transactions, Kung and Papadimitriou [89] use the metric to capture the parallelism of a locking scheme, While acknowledging that the metric is theoretical, they insist that it may have "practical significance as well, if the schedulers in question have relatively small scheduling times as compared with waiting and execution times." Herlihy [74] employed the metric to compare various optimistic and pessimistic synchronization techniques using commutativity of operations constituting high-level transactions. A synchronization technique is implicitly considered in [74] as highly concurrent, namely "optimal", if no other technique accepts more schedules. By contrast, we focus here on a *dynamic* model where the scheduler cannot use the prior knowledge of all the shared addresses to be accessed. Also, unlike [74, 89], the results in this chapter require *all* operations, including aborted ones, to observe (locally) consistent states.

Concurrency optimality. This chapter shows that "semantics-oblivious" optimistic TM and "semantics-aware" pessimistic locking are incomparable with respect to exploiting concurrency of the list-based set. Yet, we have shown how to use the benefits of optimism to derive a concurrency optimal implementation that is fine-tuned to the semantics of the list-based set. Intuitively, the ability of an implementation to successfully process interleaving steps of concurrent threads is an appealing property that should be met by performance gains. We believe this to be so.

In work that is not part of the thesis [56], we confirm experimentally that the concurrency optimal optimistic implementation of the list-based set based on I^{RM} outperforms the state-of-the-art implementations of the list-based set, namely, the *Lazy linked list* [71] and the *Harris-Michael linked list* [70, 105]. Does the claim also hold for other data structures? We suspect so. For example, similar but more general data structures, such as skip-lists or tree-based dictionaries, may allow for optimizations similar to proposed in this paper. Our results provides some preliminary hints in the quest for the "right" synchronization technique to develop highly concurrent and efficient implementations of data types.

9

Concluding remarks

Everything has to come to an end, sometime.

Lyman Frank Baum-The Marvelous Land of Oz

The inclusion of hardware support for transactions in mainstream CPU's [1, 107, 111] suggests that TM is an important concurrency abstraction. However, hardware transactions are not going to be sufficient to support efficient concurrent programming since they may be aborted spuriously; the fast but potentially unreliable hardware transactions must be complemented with slower, but more reliable software transactions. Thus, understanding the inherent cost of both hardware and software transactions is of both theoretical and practical interest.

Below, we briefly recall the outcomes of the thesis and overview the future research directions.

Safety for TMs. We formalized the semantics of a safe TM: every transaction, including aborted and incomplete ones, must observe a view that is consistent with some sequential execution. We introduced the notion of deferred-update semantics which explicitly precludes reading from a transaction that has not yet invoked tryCommit. We believe that our definition is useful to TM practitioners, since it streamlines possible implementations of t-read and tryCommit operations.

Complexity of TMs. The cost of the TM abstraction is parametrized by several properties: safety for transactions, conditions under which transactions must terminate, conditions under which transactions must commit/abort, bound on the number of versions that can be maintained and a multitude of other implementation strategies like disjoint-access parallelism and invisible reads.

At a high-level, the complexity bounds presented in the thesis suggest that providing high degrees of concurrency in software transactional memory (STM) implementations incurs a considerable synchronization cost. As we show, permissive STMs, while providing the best possible concurrency in theory, require a strong synchronization primitive (AWAR) or a memory fence (RAW) per read operation, which may result in excessively slow execution times. Progressive STMs provide only basic concurrency by adapting to data conflicts, but perform considerably better in this respect: we present progressive implementations that incur constant RAW/AWAR complexity.

Since Transactional memory was originally proposed as an alternative to locking, early STMs implementations [52, 79, 101, 117, 120] adopted optimistic concurrency control and guaranteed that a prematurely

halted transaction cannot not prevent other transactions from committing. However, popular state-of-the-art STM implementations like *TL2* [39] and *NOrec* [36] are progressive, providing no non-blocking progress guarantees for transactions, but perform empirically better than obstruction-free TMs. Complexity lower and upper bounds presented in the thesis explain this performance gap.

Do our results mean that maximizing the ability of processing multiple transactions in parallel or providing non-blocking progress should not be an important factor in STM design? It would seem so. Should we rather even focus on speculative "single-lock" solutions á la *flat combining* [72] or "pessimistic" STMs in which transactions never abort [5]? Difficult to say affirmatively, but probably not, since our results suggest progressive STMs incur low complexity overheads as also evidenced by their good empirical performance on most realistic TM workloads [36, 39].

Several questions yet remain open on the complexity of STMs. For instance, the bounds in the thesis were derived for the TM-correctness property of strict serializability and its restrictions. But there has been study of relaxations of strict serializability like *snapshot isolation* [24, 31]. Verifying if the lower bounds presented in the thesis hold under such weak TM-correctness properties and extending the proofs if indeed, presents interesting open questions. The discussion section of Chapters 4, 5 and 6 additionally list some unresolved questions closely related to the results in the thesis.

One problem of practical need that is not considered in the thesis concerns the interaction of transactional code with *non-transactional* code, *i.e.*, the same data item is accessed both transactionally and non-transactionally. It is expected that code executed within a transaction behave as lock-based code within a single "global lock" [104, 116] to avoid memory races. Techniques to ensure the safety of non-transactional accesses have been formulated through the notion of *privatization* [23, 118]. Devising techniques to ensure privatization for TMs and understanding the cost of enforcing it is an important research direction.

In the thesis, we assumed that a rmw event is an access to a single base object. However, there have been proposals to provide implementations with the ability to invoke k-rmw; $k \in \mathbb{N}$ primitives [21, 42] that allow accessing up to k base objects in a single atomic event. For example, the k-cas instruction allows to perform k cas instructions atomically on a vector $\langle b_1, \ldots, b_k \rangle$ of base objects: it accepts as input a vector $\langle old_1, \ldots, old_k, new_1, \ldots, new_k \rangle$ and atomically updates the value of $\langle b_1, \ldots, b_k \rangle$ to $\langle new_1, \ldots, new_k \rangle$ and returns *true iff* for all $i \in \{1, \ldots, k\}$, $old_i = new_i$; otherwise it returns *false*. However, the ability to access such k-rmw primitives does not necessarily simplify the design and improve the performance of non-blocking implementations nor overcome the compositionality issue [21, 42, 81]. Nonetheless, verifying if the lower bounds presented in the thesis hold in this shared memory model is an interesting problem.

HyTMs. We have introduced an analytical model for hybrid transactional memory that captures the notion of cached accesses as performed by hardware transactions. We then derived lower and upper bounds in this model to capture the inherent tradeoff between the degree of concurrency allowed among hardware and software transactions and the instrumentation overhead introduced on the hardware. In a nutshell, our results say that it is impossible to completely forgo instrumentation in a sequential HyTM, and that any opaque HyTM implementation providing non-trivial progress either has to pay a *linear* number of metadata accesses, or will have to allow slow-path transactions to *abort* fast-path operations.

Our model of HTMs assumed that the hardware resources were *bounded*, in the sense that, a hardware transaction may only access a bounded number of data items, exceeding which, it incurs a capacity abort. To overcome the inherent limitations of bounded HTMs, there have been proposals for "unbounded HTMs" that allow transactions to commit even if they exceed the hardware resources [12, 68]. The HyTM model from Chapter 7 can be easily extended to accommodate unbounded HTM designs by disregarding capacity aborts.

Some papers have investigated alternatives to providing HTMs with an STM fallback, such as *sandboxing* [4, 33], or employing *hardware-accelerated* STM [114, 119], and the use of both direct *and* cached accesses within the same hardware transaction to reduce instrumentation overhead [88, 112, 113]. Another approach proposed *reduced hardware transactions* [102], where a part of the slow-path is executed

using a short fast-path transaction, which allows to partially eliminate instrumentation from the hardware fast-path. Modelling and deriving complexity bounds for HyTM proposals outside the HyTM model described in the thesis is an interesting future direction.

Relaxed transactional memory. The concurrency lower bounds derived in Chapter 8 illustrated that a strictly serializable TM, when used as a black-box to transform a sequential implementation of the list-based set to a concurrent one, is not concurrency-optimal. This is due to the fact that TM detects conflicts at the level of transactional reads and writes resulting in *false conflicts*, in the sense that, the read-write conflict may not affect the correctness of the implemented high-level set type. As we have shown, we can derive a *concurrency optimal* optimistic (non-strictly serializable) implementation that can process every correct schedule of the list-based set. Indeed, several papers have studied "relaxed" TMs that are fined-tuned to the semantics of the high-level data type [50, 76, 77]. Exploring the complexity of such relaxed TM models represents a very important future research direction.

List of Figures

1.1 Transforming a sequential implementation of the list-based set to a TM-based concurrent one . . . 14

3.1 History H is final-state opaque, while its prefix H' is not final-state opaque. . . . 35

3.2 An infinite history in which $tryC_1$ is incomplete and any two transactions are concurrent. Each finite prefix of the history is du-opaque, but the infinite limit of the ever-extending sequence is not du-opaque. . . . 37

3.3 A history that is opaque, but not du-opaque. . . . 41

3.4 A sequential du-opaque history, which is not opaque by the definition of [59]. . . . 42

3.5 A history that is du-VWC, but not du-opaque. . . . 44

3.6 A history which is du-VWC but not du-TMS1. . . . 46

3.7 A history which is du-TMS1 but not du-VWC. . . . 46

3.8 A history that is du-opaque, but not TMS2 [43]. . . . 47

4.1 Executions in the proof of Lemma 4.1; By weak DAP, T_ϕ cannot distinguish this from the execution in Figure 4.1a . . . 51

4.2 Execution E of a permissive, opaque TM: T_2 and T_3 force T_1 to perform a RAW/AWAR in each $R_1(X_k)$, $2 \leq k \leq m$ . . . 70

5.1 Executions in the proof of Theorem 5.1; execution in 5.1d is not strictly serializable . . . 74

5.2 Executions in the proof of Theorem 5.4; execution in 5.2e is not opaque . . . 79

5.3 Complexity gap between blocking and non-blocking TMs . . . 87

6.1 Executions in the proof of Theorem 6.1; execution in 6.1a must maintain c distinct values of every t-object . . . 92

6.2 Executions in the proof of Theorem 6.2; execution in 6.2c is not strictly serializable . . . 94

6.3 Executions in the proof of Theorem 6.3; execution in 6.3c is not strictly serializable . . . 96

7.1 Tracking set aborts in fast-path transactions; we denote a fast-path (and resp. slow-path) transaction by F (and resp. S) . . . 102

7.2 Execution E in Figure 7.2a is indistinguishable to T_1 from the execution E' in Figure 7.2b 103

7.3 Executions in the proof of Theorem 7.4; execution in 7.3d is not strictly serializable . . . 104

8.1 A concurrency scenario for a list-based set, initially $\{1, 3, 4\}$, where value i is stored at node X_i: *insert*(2) and *insert*(5) can proceed concurrently with *contains*(5), the history is LS-linearizable but not serializable. (We only depict important read-write events here.) 124

8.2 (a) a history exporting schedule σ, with initial state $\{1, 2, 3\}$, accepted by $I^{LP} \in \mathcal{SM}$; (b) a history exporting a problematic schedule σ', with initial state $\{3\}$, which should be accepted by any $I \in \mathcal{P}$ if it accepts σ 126

List of Tables

3.1 Relations between TM consistency definitions. 48

4.1 Complexity bounds for progressive TMs. 71
4.2 Complexity bounds for strongly progressive TMs. 71

10 Bibliography

[1] Advanced Synchronization Facility Proposed Architectural Specification, March 2009. `http://developer.amd.com/wordpress/media/2013/09/45432-ASF_Spec_2.1.pdf`.

[2] Transactional Memory in GCC. 2012.

[3] S. V. Adve and K. Gharachorloo. Shared memory consistency models: A tutorial. *IEEE Computer*, 29(12):66–76, 1996.

[4] Y. Afek, A. Levy, and A. Morrison. Software-improved hardware lock elision. In *PODC*. ACM, 2014.

[5] Y. Afek, A. Matveev, and N. Shavit. Pessimistic software lock-elision. In *Proceedings of the 26th International Conference on Distributed Computing*, DISC'12, pages 297–311, Berlin, Heidelberg, 2012. Springer-Verlag.

[6] M. K. Aguilera, S. Frølund, V. Hadzilacos, S. L. Horn, and S. Toueg. Abortable and query-abortable objects and their efficient implementation. In *PODC*, pages 23–32, 2007.

[7] D. Alistarh, P. Eugster, M. Herlihy, A. Matveev, and N. Shavit. Stacktrack: An automated transactional approach to concurrent memory reclamation. In *Proceedings of the Ninth European Conference on Computer Systems*, EuroSys '14, pages 25:1–25:14, New York, NY, USA, 2014. ACM.

[8] D. Alistarh, J. Kopinsky, P. Kuznetsov, S. Ravi, and N. Shavit. Inherent limitations of hybrid transactional memory. *CoRR*, abs/1405.5689, 2014.

[9] D. Alistarh, J. Kopinsky, P. Kuznetsov, S. Ravi, and N. Shavit. Inherent limitations of hybrid transactional memory. *CoRR*, abs/1405.5689, 2014. To appear in 29th International Symposium on Distributed Computing (DISC'15), Japan.

[10] D. Alistarh, J. Kopinsky, P. Kuznetsov, S. Ravi, and N. Shavit. Inherent limitations of hybrid transactional memory. *6th Workshop on the Theory of Transactional Memory, Paris, France*, 2014.

[11] B. Alpern and F. B. Schneider. Defining liveness. *Inf. Process. Lett.*, 21(4):181–185, Oct. 1985.

[12] C. S. Ananian, K. Asanovic, B. C. Kuszmaul, C. E. Leiserson, and S. Lie. Unbounded transactional memory. In *Proceedings of the 11th International Symposium on High-Performance Computer Architecture*, HPCA '05, pages 316–327, Washington, DC, USA, 2005. IEEE Computer Society.

[13] J. H. Anderson and M. Moir. Universal constructions for multi-object operations. In *Proceedings of the Fourteenth Annual ACM Symposium on Principles of Distributed Computing*, PODC '95, pages 184–193, New York, NY, USA, 1995. ACM.

[14] T. E. Anderson. The performance of spin lock alternatives for shared-memory multiprocessors. *IEEE Trans. Parallel Distrib. Syst.*, 1(1):6–16, 1990.

[15] H. Attiya, A. Gotsman, S. Hans, and N. Rinetzky. Safety of live transactions in transactional memory: TMS is necessary and sufficient. In *DISC*, pages 376–390, 2014.

[16] H. Attiya, R. Guerraoui, D. Hendler, and P. Kuznetsov. The complexity of obstruction-free implementations. *J. ACM*, 56(4), 2009.

[17] H. Attiya, R. Guerraoui, D. Hendler, P. Kuznetsov, M. Michael, and M. Vechev. Laws of order: Expensive synchronization in concurrent algorithms cannot be eliminated. In *POPL*, pages 487–498, 2011.

[18] H. Attiya, S. Hans, P. Kuznetsov, and S. Ravi. Safety of deferred update in transactional memory. *2013 IEEE 33rd International Conference on Distributed Computing Systems*, 0:601–610, 2013.

[19] H. Attiya, S. Hans, P. Kuznetsov, and S. Ravi. Safety of deferred update in transactional memory. *CoRR*, abs/1301.6297, 2013.

[20] H. Attiya, S. Hans, P. Kuznetsov, and S. Ravi. Safety and deferred update in transactional memory. In R. Guerraoui and P. Romano, editors, *Transactional Memory. Foundations, Algorithms, Tools, and Applications*, volume 8913 of *Lecture Notes in Computer Science*, pages 50–71. Springer International Publishing, 2015.

[21] H. Attiya and D. Hendler. Time and space lower bounds for implementations using k-cas. *Parallel and Distributed Systems, IEEE Transactions on*, 21(2):162 –173, feb. 2010.

[22] H. Attiya, D. Hendler, and P. Woelfel. Tight rmr lower bounds for mutual exclusion and other problems. In *Proceedings of the Twenty-seventh ACM Symposium on Principles of Distributed Computing*, PODC '08, pages 447–447, New York, NY, USA, 2008. ACM.

[23] H. Attiya and E. Hillel. The cost of privatization in software transactional memory. *IEEE Trans. Computers*, 62(12):2531–2543, 2013.

[24] H. Attiya, E. Hillel, and A. Milani. Inherent limitations on disjoint-access parallel implementations of transactional memory. *Theory of Computing Systems*, 49(4):698–719, 2011.

[25] H. Attiya and A. Milani. Transactional scheduling for read-dominated workloads. In *Proceedings of the 13th International Conference on Principles of Distributed Systems*, OPODIS '09, pages 3–17, Berlin, Heidelberg, 2009. Springer-Verlag.

[26] H. Attiya, G. Ramalingam, and N. Rinetzky. Sequential verification of serializability. In *Proceedings of the 37th annual ACM SIGPLAN-SIGACT symposium on Principles of programming languages*, pages 31–42, 2010.

[27] H. Attiya and J. Welch. *Distributed Computing. Fundamentals, Simulations, and Advanced Topics.* John Wiley & Sons, 2004.

[28] G. Barnes. A method for implementing lock-free shared-data structures. In *Proceedings of the Fifth Annual ACM Symposium on Parallel Algorithms and Architectures*, SPAA '93, pages 261–270, New York, NY, USA, 1993. ACM.

[29] R. Bayer and M. Schkolnick. Concurrency of operations on B-trees. In *Readings in database systems*, pages 129–139. Morgan Kaufmann Publishers Inc., 1988.

[30] E. Bruno. What Is Priority Inversion (And How Do You Control It)? 2011.

[31] V. Bushkov, D. Dziuma, P. Fatourou, and R. Guerraoui. The pcl theorem: Transactions cannot be parallel, consistent and live. In *SPAA*, pages 178–187, 2014.

[32] V. Bushkov, R. Guerraoui, and M. Kapalka. On the liveness of transactional memory. In *Proceedings of the 2012 ACM Symposium on Principles of Distributed Computing*, PODC '12, pages 9–18, New York, NY, USA, 2012. ACM.

[33] I. Calciu, T. Shpeisman, G. Pokam, and M. Herlihy. Improved single global lock fallback for best-effort hardware transactional memory. In *Transact 2014 Workshop*. ACM, 2014.

[34] T. Crain, D. Imbs, and M. Raynal. Read invisibility, virtual world consistency and permissiveness are compatible. Research Report, ASAP - INRIA - IRISA - CNRS : UMR6074 - INRIA - Institut National des Sciences Appliquées de Rennes - Université de Rennes I, 11 2010.

[35] L. Dalessandro, F. Carouge, S. White, Y. Lev, M. Moir, M. L. Scott, and M. F. Spear. Hybrid NOrec: a case study in the effectiveness of best effort hardware transactional memory. In R. Gupta and T. C. Mowry, editors, *ASPLOS*, pages 39–52. ACM, 2011.

[36] L. Dalessandro, M. F. Spear, and M. L. Scott. Norec: Streamlining stm by abolishing ownership records. *SIGPLAN Not.*, 45(5):67–78, Jan. 2010.

[37] P. Damron, A. Fedorova, Y. Lev, V. Luchangco, M. Moir, and D. Nussbaum. Hybrid transactional memory. *SIGPLAN Not.*, 41(11):336–346, Oct. 2006.

[38] D. Dice, Y. Lev, M. Moir, and D. Nussbaum. Early experience with a commercial hardware transactional memory implementation. In *Proceedings of the 14th International Conference on Architectural Support for Programming Languages and Operating Systems*, ASPLOS XIV, pages 157–168, New York, NY, USA, 2009. ACM.

[39] D. Dice, O. Shalev, and N. Shavit. Transactional locking ii. In *Proceedings of the 20th International Conference on Distributed Computing*, DISC'06, pages 194–208, Berlin, Heidelberg, 2006. Springer-Verlag.

[40] D. Dice and N. Shavit. What really makes transactions fast? In *Transact*, 2006.

[41] E. W. Dijkstra. Solution of a problem in concurrent programming control. *Commun. ACM*, 8(9):569–, Sept. 1965.

[42] S. Doherty, D. L. Detlefs, L. Groves, C. H. Flood, V. Luchangco, P. A. Martin, M. Moir, N. Shavit, and G. L. Steele, Jr. Dcas is not a silver bullet for nonblocking algorithm design. In *Proceedings of the Sixteenth Annual ACM Symposium on Parallelism in Algorithms and Architectures*, SPAA '04, pages 216–224, New York, NY, USA, 2004. ACM.

[43] S. Doherty, L. Groves, V. Luchangco, and M. Moir. Towards formally specifying and verifying transactional memory. *Formal Asp. Comput.*, 25(5):769–799, 2013.

[44] A. Dragojević, M. Herlihy, Y. Lev, and M. Moir. On the power of hardware transactional memory to simplify memory management. In *Proceedings of the 30th Annual ACM SIGACT-SIGOPS Symposium on Principles of Distributed Computing*, PODC '11, pages 99–108, New York, NY, USA, 2011. ACM.

[45] F. Ellen, P. Fatourou, E. Kosmas, A. Milani, and C. Travers. Universal constructions that ensure disjoint-access parallelism and wait-freedom. In *PODC*, pages 115–124, 2012.

[46] F. Ellen, D. Hendler, and N. Shavit. On the inherent sequentiality of concurrent objects. *SIAM J. Comput.*, 41(3):519–536, 2012.

[47] R. Ennals. The lightweight transaction library. http://sourceforge.net/projects/libltx/files/.

[48] R. Ennals. Software transactional memory should not be obstruction-free. 2005.

[49] P. Fatourou and N. D. Kallimanis. A highly-efficient wait-free universal construction. In *Proceedings of the Twenty-third Annual ACM Symposium on Parallelism in Algorithms and Architectures*, SPAA '11, pages 325–334, New York, NY, USA, 2011. ACM.

[50] P. Felber, V. Gramoli, and R. Guerraoui. Elastic transactions. In *DISC*, pages 93–107, 2009.

[51] F. Fich, D. Hendler, and N. Shavit. On the inherent weakness of conditional synchronization primitives. In *Proceedings of the Twenty-third Annual ACM Symposium on Principles of Distributed Computing*, PODC '04, pages 80–87, New York, NY, USA, 2004. ACM.

[52] K. Fraser. Practical lock-freedom. Technical report, Cambridge University Computer Laborotory, 2003.

[53] V. Gramoli, D. Harmanci, and P. Felber. On the input acceptance of transactional memory. *Parallel Processing Letters*, 20(1):31–50, 2010.

[54] V. Gramoli, P. Kuznetsov, and S. Ravi. From sequential to concurrent: correctness and relative efficiency (ba). In *Principles of Distributed Computing (PODC)*, pages 241–242, 2012.

[55] V. Gramoli, P. Kuznetsov, and S. Ravi. Optimism for boosting concurrency. *CoRR*, abs/1203.4751, 2012.

[56] V. Gramoli, P. Kuznetsov, S. Ravi, and D. Shang. A concurrency-optimal list-based set. *CoRR*, abs/1502.01633, 2015.

[57] V. Gramoli, P. Kuznetsov, S. Ravi, and D. Shang. A concurrency-optimal list-based set (ba). *CoRR*, abs/1502.01633, 2015. To appear in 29th International Symposium on Distributed Computing (DISC'15).

[58] J. Gray and A. Reuter. *Transaction Processing: Concepts and Techniques*. Morgan Kaufmann Publishers Inc., San Francisco, CA, USA, 1st edition, 1992.

[59] R. Guerraoui, T. A. Henzinger, and V. Singh. Permissiveness in transactional memories. In *DISC*, pages 305–319, 2008.

[60] R. Guerraoui and M. Kapalka. On obstruction-free transactions. In *Proceedings of the twentieth annual symposium on Parallelism in algorithms and architectures*, SPAA '08, pages 304–313, New York, NY, USA, 2008. ACM.

[61] R. Guerraoui and M. Kapalka. On the correctness of transactional memory. In *Proceedings of the 13th ACM SIGPLAN Symposium on Principles and Practice of Parallel Programming*, PPoPP '08, pages 175–184, New York, NY, USA, 2008. ACM.

[62] R. Guerraoui and M. Kapalka. The semantics of progress in lock-based transactional memory. *SIGPLAN Not.*, 44(1):404–415, Jan. 2009.

[63] R. Guerraoui and M. Kapalka. Transactional memory: Glimmer of a theory. In *Proceedings of the 21st International Conference on Computer Aided Verification*, CAV '09, pages 1–15, Berlin, Heidelberg, 2009. Springer-Verlag.

[64] R. Guerraoui and M. Kapalka. *Principles of Transactional Memory, Synthesis Lectures on Distributed Computing Theory*. Morgan and Claypool, 2010.

[65] R. Guerraoui, M. Kapalka, and J. Vitek. Stmbench7: A benchmark for software transactional memory. *SIGOPS Oper. Syst. Rev.*, 41(3):315–324, Mar. 2007.

[66] R. Guerraoui and E. Ruppert. Linearizability is not always a safety property. In *NETYS*, pages 57–69, 2014.

[67] P. K. Hagit Attiya, Sandeep Hans and S. Ravi. What is safe in transactional memory. *4th Workshop on the Theory of Transactional Memory, Madeira, Portugal*, 2012.

[68] L. Hammond, V. Wong, M. Chen, B. D. Carlstrom, J. D. Davis, B. Hertzberg, M. K. Prabhu, H. Wijaya, C. Kozyrakis, and K. Olukotun. Transactional memory coherence and consistency. *SIGARCH Comput. Archit. News*, 32(2):102–, Mar. 2004.

[69] T. Harris, J. R. Larus, and R. Rajwar. *Transactional Memory, 2nd edition*. Synthesis Lectures on Computer Architecture. Morgan & Claypool Publishers, 2010.

[70] T. L. Harris. A pragmatic implementation of non-blocking linked-lists. In *DISC*, pages 300–314, 2001.

[71] S. Heller, M. Herlihy, V. Luchangco, M. Moir, W. N. Scherer, and N. Shavit. A lazy concurrent list-based set algorithm. In *OPODIS*, pages 3–16, 2006.

[72] D. Hendler, I. Incze, N. Shavit, and M. Tzafrir. Flat combining and the synchronization-parallelism tradeoff. In *SPAA*, pages 355–364, 2010.

[73] J. L. Hennessy and D. A. Patterson. *Computer Architecture: A Quantitative Approach*. Morgan Kaufmann Publishers Inc., San Francisco, CA, USA, 3 edition, 2003.

[74] M. Herlihy. Apologizing versus asking permission: optimistic concurrency control for abstract data types. *ACM Trans. Database Syst.*, 15(1):96–124, 1990.

[75] M. Herlihy. Wait-free synchronization. *ACM Trans. Prog. Lang. Syst.*, 13(1):123–149, 1991.

[76] M. Herlihy and E. Koskinen. Transactional boosting: A methodology for highly-concurrent transactional objects. In *PPoPP*, New York, NY, USA, 2008. ACM.

[77] M. Herlihy and E. Koskinen. Composable transactional objects: A position paper. In Z. Shao, editor, *Programming Languages and Systems*, volume 8410 of *Lecture Notes in Computer Science*, pages 1–7. Springer Berlin Heidelberg, 2014.

[78] M. Herlihy, V. Luchangco, and M. Moir. Obstruction-free synchronization: Double-ended queues as an example. In *ICDCS*, pages 522–529, 2003.

[79] M. Herlihy, V. Luchangco, M. Moir, and W. N. Scherer, III. Software transactional memory for dynamic-sized data structures. In *Proceedings of the Twenty-second Annual Symposium on Principles of Distributed Computing*, PODC '03, pages 92–101, New York, NY, USA, 2003. ACM.

[80] M. Herlihy and J. E. B. Moss. Transactional memory: architectural support for lock-free data structures. In *ISCA*, pages 289–300, 1993.

[81] M. Herlihy and N. Shavit. *The art of multiprocessor programming*. Morgan Kaufmann, 2008.

[82] M. Herlihy and N. Shavit. On the nature of progress. In *OPODIS*, pages 313–328, 2011.

[83] M. Herlihy and J. M. Wing. Linearizability: A correctness condition for concurrent objects. *ACM Trans. Program. Lang. Syst.*, 12(3):463–492, 1990.

[84] L. Hyonho. Local-spin mutual exclusion algorithms on the DSM model using fetch-and-store objects. 2003.

[85] D. Imbs and M. Raynal. Virtual world consistency: A condition for STM systems (with a versatile protocol with invisible read operations). *Theor. Comput. Sci.*, 444, July 2012.

[86] A. Israeli and L. Rappoport. Disjoint-access-parallel implementations of strong shared memory primitives. In *PODC*, pages 151–160, 1994.

[87] D. König. *Theorie der Endlichen und Unendlichen Graphen: Kombinatorische Topologie der Streckenkomplexe*. Akad. Verlag. 1936.

[88] S. Kumar, M. Chu, C. J. Hughes, P. Kundu, and A. Nguyen. Hybrid transactional memory. In *Proceedings of the Eleventh ACM SIGPLAN Symposium on Principles and Practice of Parallel Programming*, PPoPP '06, pages 209–220, New York, NY, USA, 2006. ACM.

[89] H. T. Kung and C. H. Papadimitriou. An optimality theory of concurrency control for databases. In *SIGMOD*, pages 116–126, 1979.

[90] P. Kuznetsov and S. Ravi. On the cost of concurrency in transactional memory. In *OPODIS*, pages 112–127, 2011.

[91] P. Kuznetsov and S. Ravi. On the cost of concurrency in transactional memory. *CoRR*, abs/1103.1302, 2011.

[92] P. Kuznetsov and S. Ravi. On partial wait-freedom in transactional memory. *CoRR*, abs/1407.6876, 2014.

[93] P. Kuznetsov and S. Ravi. On partial wait-freedom in transactional memory. In *Proceedings of the 2015 International Conference on Distributed Computing and Networking, ICDCN 2015, Goa, India, January 4-7, 2015*, page 10, 2015.

[94] P. Kuznetsov and S. Ravi. Progressive transactional memory in time and space. *CoRR*, abs/1502.04908, 2015.

[95] P. Kuznetsov and S. Ravi. Progressive transactional memory in time and space. *CoRR*, abs/1502.04908, 2015. To appear in 13th International Conference on Parallel Computing Technologies, Russia.

[96] P. Kuznetsov and S. Ravi. Why transactional memory should not be obstruction-free. *CoRR*, abs/1502.02725, 2015.

[97] P. Kuznetsov and S. Ravi. Why transactional memory should not be obstruction-free. *CoRR*, abs/1502.02725, 2015. To appear in 29th International Symposium on Distributed Computing (DISC'15), Japan.

[98] M. Lesani, V. Luchangco, and M. Moir. Putting opacity in its place. In *WTTM*, 2012.

[99] Y. Lev, M. Moir, and D. Nussbaum. Phtm: Phased transactional memory. In *In Workshop on Transactional Computing (Transact), 2007. research.sun.com/scalable/pubs/ TRANSACT2007PhTM.pdf.*

[100] N. A. Lynch. *Distributed Algorithms*. Morgan Kaufmann, 1996.

[101] V. J. Marathe, W. N. S. Iii, and M. L. Scott. Adaptive software transactional memory. In *In Proc. of the 19th Intl. Symp. on Distributed Computing*, pages 354–368, 2005.

[102] A. Matveev and N. Shavit. Reduced hardware transactions: a new approach to hybrid transactional memory. In *Proceedings of the 25th ACM symposium on Parallelism in algorithms and architectures*, pages 11–22. ACM, 2013.

[103] P. E. McKenney. Memory barriers: a hardware view for software hackers. Linux Technology Center, IBM Beaverton, June 2010.

[104] V. Menon, S. Balensiefer, T. Shpeisman, A.-R. Adl-Tabatabai, R. L. Hudson, B. Saha, and A. Welc. Single global lock semantics in a weakly atomic stm. *SIGPLAN Not.*, 43(5):15–26, May 2008.

[105] M. M. Michael. High performance dynamic lock-free hash tables and list-based sets. In *SPAA*, pages 73–82, 2002.

[106] M. M. Michael and M. L. Scott. Simple, fast, and practical non-blocking and blocking concurrent queue algorithms. In *PODC*, pages 267–275, 1996.

[107] M. Ohmacht. Memory Speculation of the Blue Gene/Q Compute Chip, 2011. `http://wands.cse.lehigh.edu/IBM_BQC_PACT2011.ppt`.

[108] S. S. Owicki and L. Lamport. Proving liveness properties of concurrent programs. *ACM Trans. Program. Lang. Syst.*, 4(3):455–495, 1982.

[109] C. H. Papadimitriou. The serializability of concurrent database updates. *J. ACM*, 26:631–653, 1979.

[110] D. Perelman, R. Fan, and I. Keidar. On maintaining multiple versions in STM. In *PODC*, pages 16–25, 2010.

[111] J. Reinders. Transactional Synchronization in Haswell, 2012. `http://software.intel.com/en-us/blogs/2012/02/07/transactional-synchronization-in-haswell/`.

[112] T. Riegel. Software Transactional Memory Building Blocks. 2013.

[113] T. Riegel, P. Marlier, M. Nowack, P. Felber, and C. Fetzer. Optimizing hybrid transactional memory: The importance of nonspeculative operations. In *Proceedings of the 23rd ACM Symposium on Parallelism in Algorithms and Architectures*, pages 53–64. ACM, 2011.

[114] B. Saha, A.-R. Adl-Tabatabai, and Q. Jacobson. Architectural support for software transactional memory. In *Proceedings of the 39th Annual IEEE/ACM International Symposium on Microarchitecture*, MICRO 39, pages 185–196, Washington, DC, USA, 2006. IEEE Computer Society.

[115] W. N. Scherer, III and M. L. Scott. Advanced contention management for dynamic software transactional memory. In *Proceedings of the Twenty-fourth Annual ACM Symposium on Principles of Distributed Computing*, PODC '05, pages 240–248, New York, NY, USA, 2005. ACM.

[116] M. L. Scott. *Shared-memory Synchronization, Synthesis Lectures on Distributed Computing Theory.* Morgan and Claypool, 2013.

[117] N. Shavit and D. Touitou. Software transactional memory. In *PODC*, pages 204–213, 1995.

[118] M. F. Spear, V. J. Marathe, L. Dalessandro, and M. L. Scott. Privatization techniques for software transactional memory. In *Proceedings of the Twenty-sixth Annual ACM Symposium on Principles of Distributed Computing*, PODC '07, pages 338–339, New York, NY, USA, 2007. ACM.

[119] M. F. Spear, A. Shriraman, L. Dalessandro, S. Dwarkadas, and M. L. Scott. Nonblocking transactions without indirection using alert-on-update. In *Proceedings of the Nineteenth Annual ACM Symposium on Parallel Algorithms and Architectures*, SPAA '07, pages 210–220, New York, NY, USA, 2007. ACM.

[120] F. Tabba, M. Moir, J. R. Goodman, A. W. Hay, and C. Wang. Nztm: Nonblocking zero-indirection transactional memory. In *Proceedings of the Twenty-first Annual Symposium on Parallelism in Algorithms and Architectures*, SPAA '09, pages 204–213, New York, NY, USA, 2009. ACM.

[121] G. Taubenfeld. The black-white bakery algorithm and related bounded-space, adaptive, local-spinning and fifo algorithms. In *DISC '04: Proceedings of the 23rd International Symposum on Distributed Computing*, 2004.

[122] P. K. Vincent Gramoli and S. Ravi. Sharing a sequential data structure: correctness definition and concurrency analysis. *4th Workshop on the Theory of Transactional Memory, Madeira, Portugal*, 2012.

[123] W. E. Weihl. Commutativity-based concurrency control for abstract data types. *IEEE Trans. Comput.*, 37(12):1488–1505, 1988.

[124] G. Weikum. A theoretical foundation of multi-level concurrency control. In *PODS*, pages 31–43, 1986.

[125] G. Weikum and G. Vossen. *Transactional Information Systems: Theory, Algorithms, and the Practice of Concurrency Control and Recovery.* Morgan Kaufmann, 2002.

[126] M. Yannakakis. Serializability by locking. *J. ACM*, 31(2):227–244, 1984.

Papers

The content of the thesis is based on the following tech reports and publications.

Tech reports

P. Kuznetsov and S. Ravi. Progressive transactional memory in time and space. *CoRR*, abs/1502.04908, 2015.

P. Kuznetsov and S. Ravi. Why transactional memory should not be obstruction-free. *CoRR*, abs/1502.02725, 2015.

D. Alistarh, J. Kopinsky, P. Kuznetsov, S. Ravi, and N. Shavit. Inherent limitations of hybrid transactional memory. *CoRR*, abs/1405.5689, 2014.

V. Gramoli, P. Kuznetsov, and S. Ravi. Optimism for boosting concurrency. *CoRR*, abs/1203.4751, 2012.

P. Kuznetsov and S. Ravi. On partial wait-freedom in transactional memory. *CoRR*, abs/1407.6876, 2014.

H. Attiya, S. Hans, P. Kuznetsov, and S. Ravi. Safety of deferred update in transactional memory. *CoRR*, abs/1301.6297, 2013.

P. Kuznetsov and S. Ravi. On the cost of concurrency in transactional memory. *CoRR*, abs/1103.1302, 2011.

Publications

P. Kuznetsov and S. Ravi. Progressive transactional memory in time and space. *CoRR*, abs/1502.04908, 2015. To appear in 13th International Conference on Parallel Computing Technologies, Russia.

P. Kuznetsov and S. Ravi. Why transactional memory should not be obstruction-free. *CoRR*, abs/1502.02725, 2015. To appear in 29th International Symposium on Distributed Computing (DISC'15), Japan.

D. Alistarh, J. Kopinsky, P. Kuznetsov, S. Ravi, and N. Shavit. Inherent limitations of hybrid transactional memory. *CoRR*, abs/1405.5689, 2014. To appear in 29th International Symposium on Distributed Computing (DISC'15), Japan.

H. Attiya, S. Hans, P. Kuznetsov, and S. Ravi. Safety and deferred update in transactional memory. In R. Guerraoui and P. Romano, editors, *Transactional Memory. Foundations, Algorithms, Tools, and Applications*, volume 8913 of *Lecture Notes in Computer Science*, pages 50–71. Springer International Publishing, 2015.

P. Kuznetsov and S. Ravi. On partial wait-freedom in transactional memory. In *Proceedings of the 2015 International Conference on Distributed Computing and Networking, ICDCN 2015, Goa, India, January 4-7, 2015*, page 10, 2015.

H. Attiya, S. Hans, P. Kuznetsov, and S. Ravi. Safety of deferred update in transactional memory. *2013 IEEE 33rd International Conference on Distributed Computing Systems*, 0:601–610, 2013.

V. Gramoli, P. Kuznetsov, and S. Ravi. From sequential to concurrent: correctness and relative efficiency (ba). In *Principles of Distributed Computing (PODC)*, pages 241–242, 2012.

P. Kuznetsov and S. Ravi. On the cost of concurrency in transactional memory. In *OPODIS*, pages 112–127, 2011.

Workshop papers

D. Alistarh, J. Kopinsky, P. Kuznetsov, S. Ravi, and N. Shavit. Inherent limitations of hybrid transactional memory. *6th Workshop on the Theory of Transactional Memory, Paris, France*, 2014.

P. K. Vincent Gramoli and S. Ravi. Sharing a sequential data structure: correctness definition and concurrency analysis. *4th Workshop on the Theory of Transactional Memory, Madeira, Portugal*, 2012.

P. K. Hagit Attiya, Sandeep Hans and S. Ravi. What is safe in transactional memory. *4th Workshop on the Theory of Transactional Memory, Madeira, Portugal*, 2012.

Concurrently, I was also involved in the following paper whose contents are not included in the thesis.

V. Gramoli, P. Kuznetsov, S. Ravi, and D. Shang. A concurrency-optimal list-based set. *CoRR*, abs/1502.01633, 2015.

V. Gramoli, P. Kuznetsov, S. Ravi, and D. Shang. A concurrency-optimal list-based set (ba). *CoRR*, abs/1502.01633, 2015. To appear in 29th International Symposium on Distributed Computing (DISC'15).